


# AltiVec™ Technology Programming Environments Manual

PRELIMINARY

© Motorola Inc, 1998. All rights reserved.  
Portions hereof © International Business Machines Corp. 1991–1998. All rights reserved.

This document contains information on a new product under development. Motorola reserves the right to change or discontinue this product without notice. Information in this document is provided solely to enable system and software implementers to use PowerPC microprocessors. There are no express or implied copyright licenses granted hereunder to design or fabricate PowerPC integrated circuits or integrated circuits based on the information in this document.

Motorola reserves the right to make changes without further notice to any products herein. Motorola makes no warranty, representation or guarantee regarding the suitability of its products for any particular purpose, nor does Motorola assume any liability arising out of the application or use of any product or circuit, and specifically disclaims any and all liability, including without limitation consequential or incidental damages. "Typical" parameters can and do vary in different applications. All operating parameters, including "Typicals" must be validated for each customer application by customer's technical experts. Motorola does not convey any license under its patent rights nor the rights of others. Motorola products are not designed, intended, or authorized for use as components in systems intended for surgical implant into the body, or other applications intended to support or sustain life, or for any other application in which the failure of the Motorola product could create a situation where personal injury or death may occur. Should Buyer purchase or use Motorola products for any such unintended or unauthorized application, Buyer shall indemnify and hold Motorola and its officers, employees, subsidiaries, affiliates, and distributors harmless against all claims, costs, damages, and expenses, and reasonable attorney fees arising out of, directly or indirectly, any claim of personal injury or death associated with such unintended or unauthorized use, even if such claim alleges that Motorola was negligent regarding the design or manufacture of the part.

Motorola and  are registered trademarks and Altivec is a trademark of Motorola, Inc. Motorola, Inc. is an Equal Opportunity/Affirmative Action Employer.

The PowerPC name and the PowerPC logotype are trademarks of International Business Machines Corporation used by Motorola under license from International Business Machines Corporation.

PRELIMINARY

Overview	1
Altivec Register Set	2
Operand Conventions	3
Addressing Modes and Instruction Set Summary	4
Cache, Exceptions, and Memory Management	5
Altivec Instructions	6
Altivec Instruction Set Listings	A
Glossary of Terms and Abbreviations	GLO
Index	IND

1

Overview

2

Altivec Register Set

3

Operand Conventions

4

Addressing Modes and Instruction Set Summary

5

Cache, Exceptions, and Memory Management

6

Altivec Instructions

A

Altivec Instruction Set Listings

GLO

Glossary of Terms and Abbreviations

IND

Index

PRELIMINARY

# CONTENTS

Paragraph Number	Title	Page Number
---------------------	-------	----------------

## About This Book

Audience .....	xx
Organization.....	xx
Suggested Reading.....	xxi
General Information.....	xxi
PowerPC Documentation.....	xxii
Conventions .....	xxiii
Acronyms and Abbreviations .....	xxiv
Terminology Conventions .....	xxvi

## Chapter 1 Overview

1.1	Overview.....	1-2
1.1.1	The 64-Bit AltiVec Technology and the 32-Bit Subset.....	1-4
1.1.2	The Levels of the AltiVec ISA .....	1-5
1.1.3	Features Not Defined by the AltiVec ISA .....	1-5
1.2	The AltiVec Architectural Model .....	1-6
1.2.1	AltiVec Registers and Programming Model.....	1-6
1.2.2	Operand Conventions .....	1-7
1.2.2.1	Byte Ordering .....	1-7
1.2.2.2	Floating-Point Conventions .....	1-8
1.2.3	AltiVec Instruction Set and Addressing Modes .....	1-8
1.2.3.1	AltiVec Instruction Set .....	1-10
1.2.4	AltiVec Cache Model .....	1-11
1.2.5	AltiVec Exception Model .....	1-11
1.2.6	Memory Management Model .....	1-11

## Chapter 2 AltiVec Register Set

2.1	AltiVec Register File .....	2-1
2.1.1	The Vector Status and Control Register (VSCR) .....	2-2
2.1.2	VRSAVE Register (VRSAVE) .....	2-4
2.1.3	PowerPC Condition Register .....	2-5
2.1.4	AltiVec Bit in the PowerPC Machine State Register (MSR) .....	2-6
2.1.5	Machine Status Save/Restore Registers (SRR) .....	2-8
2.1.5.1	Machine Status Save/Restore Register 0 (SRR0).....	2-8
2.1.6	Machine Status Save/Restore Register 1 (SRR1).....	2-8
2.2	PowerPC Register Set .....	2-9

# CONTENTS

Paragraph Number	Title	Page Number
<b>Chapter 3</b>		
<b>Operand Conventions</b>		
3.1	Data Organization in Memory and Data Transfers .....	3-1
3.1.1	Aligned and Misaligned Accesses .....	3-1
3.1.2	AltiVec Byte Ordering.....	3-2
3.1.2.1	Big-Endian Byte Ordering .....	3-3
3.1.2.2	Little-Endian Byte Ordering .....	3-3
3.1.3	Quad Word Byte Ordering Example .....	3-3
3.1.4	Aligned Scalars in Little-Endian Mode .....	3-5
3.1.5	Vector Register and Memory Access Alignment .....	3-7
3.1.6	Quad-Word Data Alignment.....	3-7
3.1.6.1	Accessing a Misaligned Quad Word in Big-Endian Mode .....	3-8
3.1.6.2	Accessing a Misaligned Quad Word in Little-Endian Mode.....	3-10
3.1.6.3	Scalar Loads and Stores.....	3-11
3.1.6.3.1	Misaligned Scalar Loads and Stores.....	3-12
3.1.7	Mixed-Endian Systems .....	3-12
3.2	AltiVec Floating-Point —UISA .....	3-12
3.2.1	Floating-Point Modes .....	3-13
3.2.1.1	Java Mode .....	3-13
3.2.1.2	Non-Java Mode.....	3-14
3.2.2	Floating-Point Infinities.....	3-14
3.2.3	Floating-Point Rounding.....	3-14
3.2.4	Floating-Point Exceptions.....	3-14
3.2.4.1	NaN Operand Exception.....	3-15
3.2.4.2	Invalid Operation Exception.....	3-16
3.2.4.3	Zero Divide Exception.....	3-16
3.2.4.4	Log of Zero Exception.....	3-16
3.2.4.5	Overflow Exception.....	3-16
3.2.4.6	Underflow Exception.....	3-17
3.2.5	Floating-Point NaNs .....	3-17
3.2.5.1	NaN Precedence.....	3-17
3.2.5.2	SNaN Arithmetic .....	3-18
3.2.5.3	QNaN Arithmetic.....	3-18
3.2.5.4	NaN Conversion to Integer .....	3-18
3.2.5.5	NaN Production .....	3-18

## **Chapter 4**

### **Addressing Modes and Instruction Set Summary**

4.1	Conventions .....	4-2
4.1.1	Execution Model.....	4-2
4.1.2	Computation Modes.....	4-2

# CONTENTS

Paragraph Number	Title	Page Number
4.1.3	Classes of Instructions .....	4-2
4.1.4	Memory Addressing .....	4-3
4.1.4.1	Memory Operands .....	4-3
4.1.4.2	Effective Address Calculation .....	4-3
4.2	AltiVec UISA Instructions.....	4-4
4.2.1	Vector Integer Instructions .....	4-4
4.2.1.1	Saturation Detection .....	4-4
4.2.1.2	Vector Integer Arithmetic Instructions .....	4-5
4.2.1.3	Vector Integer Compare Instructions.....	4-12
4.2.1.4	Vector Integer Logical Instructions .....	4-14
4.2.1.5	Vector Integer Rotate and Shift Instructions .....	4-15
4.2.2	Vector Floating-Point Instructions.....	4-16
4.2.2.1	Floating-Point Division and Square-Root .....	4-16
4.2.2.1.1	Floating-Point Division .....	4-17
4.2.2.1.2	Floating-Point Square-Root .....	4-17
4.2.2.2	Floating-Point Arithmetic Instructions .....	4-18
4.2.2.3	Floating-Point Multiply-Add Instructions .....	4-18
4.2.2.4	Floating-Point Rounding and Conversion Instructions .....	4-19
4.2.2.5	Floating-Point Compare Instructions.....	4-20
4.2.2.5.1	Unordered Compares.....	4-20
4.2.2.6	Floating-Point Estimate Instructions .....	4-23
4.2.3	Load and Store Instructions .....	4-23
4.2.3.1	Alignment .....	4-24
4.2.3.2	Load and Store Address Generation .....	4-24
4.2.3.2.1	Register Indirect with Index Addressing for Loads and Stores.....	4-24
4.2.3.3	Vector Load Instructions .....	4-25
4.2.3.3.1	Vector Load Instructions Supporting Alignment .....	4-26
4.2.3.4	Vector Store Instructions .....	4-28
4.2.4	Control Flow .....	4-28
4.2.5	Vector Permutation and Formatting Instructions.....	4-28
4.2.5.1	Vector Pack Instructions.....	4-29
4.2.5.2	Vector Unpack Instructions .....	4-30
4.2.5.3	Vector Merge Instructions .....	4-31
4.2.5.4	Vector Splat Instructions .....	4-32
4.2.5.5	Vector Permute Instructions .....	4-33
4.2.5.6	Vector Select Instruction .....	4-33
4.2.5.7	Vector Shift Instructions.....	4-34
4.2.5.7.1	Immediate Interelement Shifts/Rotates.....	4-34
4.2.5.7.2	Computed Interelement Shifts/Rotates .....	4-35
4.2.5.7.3	Variable Interelement Shifts .....	4-35
4.2.6	Processor Control Instructions—UISA .....	4-36
4.2.6.1	AltiVec Status and Control Register Instructions.....	4-36
4.2.7	Recommended Simplified Mnemonics.....	4-37

# CONTENTS

Paragraph Number	Title	Page Number
4.3	AltiVec VEA Instructions .....	4-38
4.3.1	Memory Control Instructions—VEA .....	4-38
4.3.1.1	User-Level Cache Instructions—VEA .....	4-38

## Chapter 5 Cache, Exceptions, and Memory Management

5.1	Memory Bandwidth Management .....	5-1
5.1.1	Software-Directed Prefetch.....	5-1
5.1.1.1	Data Stream Touch (dst).....	5-1
5.1.1.2	Transient Streams (dstt).....	5-3
5.1.1.3	Storing to Streams (dstst).....	5-4
5.1.1.4	Stopping Streams (dss) .....	5-4
5.1.1.5	Exception Behavior of Prefetch Streams .....	5-5
5.1.1.6	Synchronization Behavior of Streams .....	5-6
5.1.1.7	Address Translation for Streams .....	5-6
5.1.1.8	Stream Usage Notes.....	5-7
5.1.1.9	Stream Implementation Assumptions.....	5-8
5.1.2	Prioritizing Cache Block Replacement.....	5-9
5.1.2.1	Partially Executed AltiVec Instructions .....	5-9
5.2	DSI Exception—Data Address Breakpoint .....	5-9
5.3	AltiVec Unavailable Exception (0x00F20) .....	5-9

## Chapter 6 AltiVec Instructions

6.1	Instruction Formats .....	6-1
6.1.1	Instruction Fields .....	6-2
6.1.2	Notation and Conventions .....	6-2
6.2	AltiVec Instruction Set .....	6-9

## Appendix A AltiVec Instruction Set Listings

A.1	Instructions Sorted by Mnemonic.....	A-1
A.2	Instructions Sorted by Opcode.....	A-7
A.3	Instructions Sorted by Form.....	A-13
A.4	Instruction Set Legend .....	A-18



# ILLUSTRATIONS

Figure Number	Title	Page Number
Figure 1-1	High Level Structural Overview of PowerPC with AltiVec Technology.....	1-3
Figure 1-2	AltiVec Top-Level Diagram .....	1-6
Figure 1-3	Big-Endian Byte Ordering for a Vector Register.....	1-7
Figure 1-4	Intraelement Example, vaddsws.....	1-8
Figure 1-5	Interelement Example, vperm .....	1-9
Figure 2-1	AltiVec Register File.....	2-2
Figure 2-2	Vector Status and Control Register (VSCR).....	2-2
Figure 2-3	VSCR Moved to a Vector Register .....	2-3
Figure 2-4	Saving/Restoring the AltiVec Context Register (VRSAVE).....	2-4
Figure 2-5	Condition Register (CR) .....	2-5
Figure 2-6	Machine State Register (MSR)—64-Bit Implementation.....	2-6
Figure 2-7	Machine State Register (MSR)—32-Bit Implementation.....	2-6
Figure 2-8	Machine Status Save/Restore Register 0 (SRR0) .....	2-8
Figure 2-9	Machine Status Save/Restore Register 1 (SRR1) .....	2-8
Figure 2-10	OEA Programming Model—All Registers .....	2-10
Figure 3-1	Big-Endian Mapping of a Quad Word.....	3-4
Figure 3-2	Little-Endian Mapping of a Quad Word.....	3-4
Figure 3-3	Little-Endian Mapping of Quad Word—Alternate View .....	3-4
Figure 3-4	Quad Word Load with Munged Little-Endian Applied .....	3-5
Figure 3-5	AltiVec Little Endian Swap .....	3-6
Figure 3-6	Misaligned Vector in Big-Endian Addressing Mode.....	3-8
Figure 3-7	Misaligned Vector in Little-Endian Addressing Mode.....	3-8
Figure 3-8	Big-Endian Quad Word Alignment .....	3-9
Figure 3-9	Little-Endian Alignment .....	3-11
Figure 4-1	Register Indirect with Index Addressing for Loads/Stores .....	4-25
Figure 5-1	Format of rB in dst Instruction.....	5-2
Figure 5-2	Data Stream Touch.....	5-3
Figure 6-1	. Data Stream Touch .....	6-10
Figure 6-2	. Data Stream Touch .....	6-13
Figure 6-3	. Data Stream Touch .....	6-14
Figure 6-4	Data Stream Touch.....	6-15
Figure 6-5	.....	6-16
Figure 6-6	. Data Stream Touch .....	6-16
Figure 6-7	Load Instructions.....	6-19
Figure 6-8	Load Vector for Shift Left.....	6-22
Figure 6-9	Instruction vperm Used in Aligning Data .....	6-23
Figure 6-10	Store Instructions .....	6-30
Figure 6-11	Basic 2-Source Operands —32-Bit Elements .....	6-37
Figure 6-12	Basic 2-Source Operands, Sixteen 8-Bit Elements.....	6-38
Figure 6-13	Basic 2-Source Operands, Eight 16-Bit Elements .....	6-39
Figure 6-14	Basic 2-Source Operands, Four 32-bit Elements .....	6-41
Figure 6-15	Basic 2-Source Operands, Sixteen 8-bit Elements.....	6-42
Figure 6-16	Basic 2-Source Operands, Sixteen 8-bit Elements.....	6-43

# ILLUSTRATIONS

<b>Figure Number</b>	<b>Title</b>	<b>Page Number</b>
Figure 6-17	Basic 2-Source Operands, Eight 16-bit Elements .....	6-44
Figure 6-18	Basic 2-Source Operands, Eight 16-bit Elements .....	6-45
Figure 6-19	Basic 2-Source Operands, Four 32-bit Elements .....	6-46
Figure 6-20	Basic 2-Source Operands, Four 32-bit Elements .....	6-47
Figure 6-21	Basic 2-Source Operands .....	6-48
Figure 6-22	Basic 2-Source Operands .....	6-49
Figure 6-23	Basic 2-Source Operands—Sixteen 8-bit Elements.....	6-50
Figure 6-24	Basic 2-Source Operands, Eight 16-bit Elements .....	6-51
Figure 6-25	Basic 2-Source Operands, Four 32-bit Elements .....	6-52
Figure 6-26	Basic 2-Source Operands—Sixteen 8-bit Elements.....	6-53
Figure 6-27	Basic 2-Source Operands, Eight 16-bit Elements .....	6-54
Figure 6-28	Basic 2-Source Operands, Four 32-bit Elements .....	6-55
Figure 6-29	Convert Four 32-bit Integer Elements to Four 32-bit Floating Point Elements.....	6-56
Figure 6-30	Convert Four 32-bit Integer Elements to Four 32-bit Floating Point Elements.....	6-57
Figure 6-31	Basic 2-Source Operands , 32-Bit Elements .....	6-59
Figure 6-32	Basic 2-Source Operands, Four 32-bit Elements .....	6-60
Figure 6-33	Basic 2-Source Operands—Sixteen 8-bit Elements.....	6-61
Figure 6-34	Basic 2-Source Operands, Eight 16-bit Elements .....	6-63
Figure 6-35	Basic 2-Source Operands, Four 32-bit Elements .....	6-64
Figure 6-36	Basic 2-Source Operands, Four 32-bit Elements .....	6-66
Figure 6-37	Basic 2-Source Operands—Sixteen 8-bit Elements.....	6-68
Figure 6-38	Basic 2-Source Operands, Eight 16-bit Elements .....	6-70
Figure 6-39	Basic 2-Source Operands, Four 32-bit Elements .....	6-71
Figure 6-40	Basic 2-Source Operands—Sixteen 8-bit Elements.....	6-72
Figure 6-41	Basic 2-Source Operands, Eight 16-bit Elements .....	6-74
Figure 6-42	Basic 2-Source Operands, Four 32-bit Elements .....	6-75
Figure 6-43	Convert Floating Point to Integer.....	6-76
Figure 6-44	Convert Floating Point to Integer.....	6-77
Figure 6-45	Basic One-Source Operands, 32-Bit Elements .....	6-79
Figure 6-46	Basic One-Source Operands, 32-Bit Elements .....	6-81
Figure 6-47	Multiply Add Floating Point .....	6-82
Figure 6-48	Basic 2-Source Operands—Four 32-bit Elements .....	6-84
Figure 6-49	Basic 2-Source Operands , 8-Bit Elements .....	6-85
Figure 6-50	Basic 2-Source Operands, Eight 16-bit Elements .....	6-86
Figure 6-51	Basic 2-Source Operands—Four 32-bit Elements .....	6-87
Figure 6-52	Basic 2-Source Operands , 8-Bit Elements .....	6-88
Figure 6-53	Basic 2-Source Operands, Eight 16-bit Elements .....	6-89
Figure 6-54	Basic 2-Source Operands—Four 32-bit Elements .....	6-90
Figure 6-55	Multiply-High and Add.....	6-91
Figure 6-56	Multiply-High Round and Add .....	6-92
Figure 6-57	Basic 2-Source Operands —32-Bit Elements .....	6-93
Figure 6-58	Basic 2-Source Operands , 8-Bit Elements .....	6-94
Figure 6-59	Basic 2-Source Operands , 16-Bit Elements .....	6-95

# ILLUSTRATIONS

Figure Number	Title	Page Number
Figure 6-60	Basic 2-Source Operands , 32-Bit Elements .....	6-96
Figure 6-61	Basic 2-Source Operands , 8-bit Elements.....	6-97
Figure 6-62	Basic 2-Source Operands , 16-Bit Elements .....	6-98
Figure 6-63	Basic 2-Source Operands , 32-Bit Elements .....	6-99
Figure 6-64	Multiply-Low and Add .....	6-100
Figure 6-65	Merge High—8-Bit Elements .....	6-101
Figure 6-66	Merge High—16-Bit Elements .....	6-102
Figure 6-67	Merge High—32-Bit Elements .....	6-103
Figure 6-68	Merge Low—8-Bit Elements .....	6-104
Figure 6-69	Merge Low—16-Bit Elements .....	6-105
Figure 6-70	Merge Low—32-Bit Elements .....	6-106
Figure 6-71	Multiply-Sum—8-Bit Elements .....	6-107
Figure 6-72	Multiply-Sum—16-Bit Elements .....	6-108
Figure 6-73	Multiply-Sum—16-Bit Elements .....	6-109
Figure 6-74	Multiply-Sum—8-Bit Elements .....	6-110
Figure 6-75	Multiply-Sum, 16-Bit Elements .....	6-112
Figure 6-76	Multiply-Sum—16-Bit Elements .....	6-113
Figure 6-77	Multiply, Even Elements, Full-Product, 8-Bit Elements .....	6-114
Figure 6-78	Multiply, Even Elements, Full-Product, 16-Bit Elements .....	6-115
Figure 6-79	Multiply, Even Elements, Full-Product, 8-Bit Elements .....	6-116
Figure 6-80	Multiply, Even Elements, Full-Product, 16-Bit Elements .....	6-117
Figure 6-81	Multiply, Odd Elements, Full-Product, 8-Bit Elements.....	6-118
Figure 6-82	Multiply, Odd Elements, Full-Product, 16-Bit Elements.....	6-119
Figure 6-83	Multiply, Odd Elements, Full-Product, 8-Bit Elements.....	6-120
Figure 6-84	Multiply, Odd Elements, Full-Product, 16-Bit Elements.....	6-121
Figure 6-85	Basic 2-Source Operands —8-Bit Elements .....	6-123
Figure 6-86	Basic 2-Source Operands —8-Bit Elements .....	6-124
Figure 6-87	Permute-8-Bit Elements .....	6-125
Figure 6-88	Pack, 32-Bit Pixels to 1/5/5/5.....	6-126
Figure 6-89	Pack, 16-Bit Elements.....	6-127
Figure 6-90	Pack—16-Bit Elements.....	6-128
Figure 6-91	Pack—32-Bit Elements .....	6-129
Figure 6-92	Pack—32-Bit Elements .....	6-130
Figure 6-93	Pack—16-Bit Elements.....	6-131
Figure 6-94	Pack—16-Bit Elements .....	6-132
Figure 6-95	Pack—32-Bit Elements .....	6-133
Figure 6-96	Pack—32-Bit Elements.....	6-134
Figure 6-97	Basic One-Source Operands—32-Bit Elements .....	6-135
Figure 6-98	Basic One-Source Operands—32-Bit Elements .....	6-136
Figure 6-99	Basic One-Source Operands—32-Bit Elements .....	6-137
Figure 6-100	Basic One-Source Operands—32-Bit Elements .....	6-138
Figure 6-101	Basic One-Source Operands—32-Bit Elements .....	6-139
Figure 6-102	Rotates and Shifts—8-Bit Elements.....	6-140

# ILLUSTRATIONS

Figure Number	Title	Page Number
Figure 6-103	Rotates and Shifts—16-Bit Elements.....	6-141
Figure 6-104	Rotates and Shifts—32-Bit Elements.....	6-142
Figure 6-105	Basic One-Source Operands, 32-Bit Elements .....	6-143
Figure 6-106	Select—1-Bit Elements .....	6-144
Figure 6-107	Vector Shift Left (Shows Shift Count = 6) .....	6-145
Figure 6-108	Rotates and Shifts—8-Bit Elements.....	6-146
Figure 6-109	Shift Left Double (Shows Shift Count = 4) .....	6-147
Figure 6-110	Rotates and Shifts, 16-Bit Elements.....	6-148
Figure 6-111	Shift Left Element (Shows Shift Count = 4).....	6-149
Figure 6-112	Rotates and Shifts, 32-Bit Elements.....	6-150
Figure 6-113	Splat 8-Bit Element from vB[Element 7] (UIMM=7) .....	6-151
Figure 6-114	Splat 16-Bit Contents from vB[Element 1] (UIMM = 1) .....	6-152
Figure 6-115	Splat 8-Bit Value from SIMM .....	6-153
Figure 6-116	Splat 16-Bit Value from SIMM .....	6-154
Figure 6-117	Splat 32-bit value from SIMM.....	6-155
Figure 6-118	Splat 8-Bit Contents from vB[Element 7] (UIMM = 7) .....	6-156
Figure 6-119	Vector Shift Right (Shift Count = 5 Shown).....	6-158
Figure 6-120	Rotates and Shifts—8-Bit Elements.....	6-159
Figure 6-121	Rotates and Shifts—16-Bit Elements.....	6-160
Figure 6-122	Rotates and Shifts—32-Bit Elements.....	6-161
Figure 6-123	Rotates and Shifts—8-Bit Elements.....	6-162
Figure 6-124	Rotates and Shifts—16-Bit Elements.....	6-163
Figure 6-125	Shift Right Element (Shift Count = 5 Shown) .....	6-164
Figure 6-126	Rotates and Shifts—8-Bit Elements.....	6-165
Figure 6-127	Basic 2-Source Operands —32-Bit Elements .....	6-167
Figure 6-128	Basic 2-Source Operands , 8-Bit Elements .....	6-168
Figure 6-129	Basic 2-Source Operands , 16-Bit Elements .....	6-169
Figure 6-130	Basic 2-Source Operands , 32-Bit Elements .....	6-170
Figure 6-131	Basic 2-Source Operands , 8-Bit Elements .....	6-171
Figure 6-132	Basic 2-Source Operands , 8-Bit Elements .....	6-172
Figure 6-133	Basic 2-Source Operands , 16-Bit Elements .....	6-173
Figure 6-134	Basic 2-Source Operands , 16-Bit Elements .....	6-174
Figure 6-135	Basic 2-Source Operands , 32-Bit Elements .....	6-175
Figure 6-136	Basic 2-Source Operands , 32-Bit Elements .....	6-176
Figure 6-137	Sum-Across, 32-Bit Elements .....	6-177
Figure 6-138	Partial (1/2) Sum Across, 32-Bit Elements .....	6-178
Figure 6-139	Partial (1/4) Sum-Across, 8-Bit Elements.....	6-179
Figure 6-140	Partial (1/4) Sum-Across, 8-Bit Elements.....	6-180
Figure 6-141	Partial (1/4) Sum-Across, 8-Bit Elements.....	6-181
Figure 6-142	Unpack High, 16-Bit Pixels .....	6-182
Figure 6-143	Unpack High, 8-Bit Elements .....	6-184
Figure 6-144	Unpack High, 16-Bit Elements .....	6-185
Figure 6-145	Unpack Low—16-Bit Pixels .....	6-186

# ILLUSTRATIONS

Figure Number	Title	Page Number
Figure 6-146	Unpack Low—8-Bit Elements .....	6-187
Figure 6-147	Unpack Low—16-Bit Elements .....	6-188

PRELIMINARY

# ILLUSTRATIONS

Figure  
Number

Title

Page  
Number

**PRELIMINARY**

# TABLES

Table Number	Title	Page Number
Table i	Acronyms and Abbreviated Terms.....	xxiv
Table ii	Terminology Conventions .....	xxvi
Table iii	Instruction Field Conventions .....	xxvii
Table 2-1	VSCR Field Descriptions .....	2-3
Table 2-2	VRSAVE Bit Settings .....	2-5
Table 2-3	CR6 Field Bit Settings for Vector Compare Instructions.....	2-6
Table 2-4	MSR Bit Settings Affected by AltiVec .....	2-7
Table 3-1	Memory Operand Alignment .....	3-2
Table 3-2	Effective Address Modifications .....	3-5
Table 4-1	Vector Integer Arithmetic Instructions.....	4-5
Table 4-2	CR6 Field Bit Settings for Vector Integer Compare Instructions .....	4-13
Table 4-3	Vector Integer Compare Instructions .....	4-13
Table 4-4	Vector Integer Logical Instructions.....	4-14
Table 4-6	Vector Integer Shift Instructions .....	4-15
Table 4-5	Vector Integer Rotate Instructions.....	4-15
Table 4-7	Floating-Point Arithmetic Instructions.....	4-18
Table 4-8	Floating-Point Multiply-Add Instructions .....	4-19
Table 4-9	Floating-Point Rounding and Conversion Instructions .....	4-20
Table 4-10	Common Mathematical Predicates.....	4-21
Table 4-11	Other Useful Predicates .....	4-21
Table 4-12	Floating-Point Compare Instructions .....	4-22
Table 4-13	Floating-Point Estimate Instructions .....	4-23
Table 4-14	Effective Address Alignment .....	4-25
Table 4-15	Integer Load Instructions.....	4-26
Table 4-16	Vector Load Instructions Supporting Alignment .....	4-27
Table 4-17	Integer Store Instructions .....	4-28
Table 4-18	Vector Pack Instructions.....	4-29
Table 4-19	Vector Unpack Instructions.....	4-31
Table 4-20	Vector Merge Instructions .....	4-32
Table 4-21	Vector Splat Instructions .....	4-33
Table 4-22	Vector Permute Instruction.....	4-33
Table 4-23	Vector Select Instruction .....	4-34
Table 4-24	Vector Shift Instructions.....	4-34
Table 4-25	Coding Various Shifts and Rotates with the vsldoi Instruction .....	4-35
Table 4-26	Move to/from Condition Register Instructions.....	4-36
Table 4-27	Simplified Mnemonics for Data Stream Touch (dst) .....	4-37
Table 4-28	User-Level Cache Instructions .....	4-39
Table 5-1	AltiVec Unavailable Exception—Register Settings.....	5-10
Table 5-2	Exception Priorities (Synchronous/Precise Exceptions) .....	5-11
Table 6-1	Instruction Syntax Conventions .....	6-2
Table 6-2	Notation and Conventions .....	6-2
Table 6-3	Instruction Field Conventions .....	6-8
Table 6-4	Precedence Rules.....	6-8

# TABLES

Table Number	Title	Page Number
Table 6-5	Instruction Description .....	6-9
Table A-1	Complete Instruction List Sorted by Mnemonic .....	A-1
Table A-2	Instructions Sorted by Opcode .....	A-7
Table A-3	Loads and Stores (opcode 31) .....	A-13
Table C-1	VA-Form .....	A-13
Table C-2	VX-Form .....	A-14
Table A-4	AltiVec Instruction Set Legend .....	A-18

PRELIMINARY



# About This Book

---

The primary objective of this manual is to help programmers provide software that is compatible across the family of PowerPC™ processors using AltiVec™ technology.

This book describes how the AltiVec technology relates to both the 64- and the 32-bit portions of the PowerPC architecture.

To locate any published errata or updates for this document, refer to the website at <http://www.mot.com/SPS/PowerPC/>.

*AltiVec Technology Programming Environments Manual* (AltiVec PEM) is used as a reference guide for programmers. The AltiVec PEM provides a description for each instruction that includes the instruction format, an individualized legend that provides such information as the level(s) of the PowerPC architecture in which the instruction may be found, the privilege level of the instruction, and figures to help in understanding how the instruction works.

Because it is important to distinguish between the levels of the PowerPC architecture in order to ensure compatibility across multiple platforms, those distinctions are shown clearly throughout this book. Most the discussions on the AltiVec technology are at the UISA level. The level of the architecture to which text refers is indicated in the outer margin, using the conventions shown in “Conventions,” on page xxiii.

This document stays consistent with the PowerPC architecture in referring to three levels, or programming environments, which are as follows:

- U** • PowerPC user instruction set architecture (UISA)—The UISA defines the level of the architecture to which user-level software should conform. The UISA defines the base user-level instruction set, user-level registers, data types, memory conventions, and the memory and programming models seen by application programmers.
- V** • PowerPC virtual environment architecture (VEA)—The VEA, which is the smallest component of the PowerPC architecture, defines additional user-level functionality that falls outside typical user-level software requirements. The VEA describes the memory model for an environment in which multiple processors or other devices can access external memory, and defines aspects of the cache model and cache control instructions from a user-level perspective. The resources defined by the VEA are particularly useful for optimizing memory accesses and for managing resources in an environment in which other processors and other devices can access external memory.

Implementations that conform to the PowerPC VEA also adhere to the UISA, but may not necessarily adhere to the OEA.

- PowerPC operating environment architecture (OEA)—The OEA defines supervisor-level resources typically required by an operating system. The OEA defines the PowerPC memory management model, supervisor-level registers, and the exception model.

Implementations that conform to the PowerPC OEA also conform to the PowerPC UISA and VEA.

For ease in reference, this book and the processor user's manuals have arranged the architecture information into topics that build upon one another, beginning with a description and complete summary of registers and instructions (for all three environments) and progressing to more specialized topics such as the cache, exception, and memory management models. As such, chapters may include information from multiple levels of the architecture but when discussing OEA and VEA, this will be noted in the text.

It is beyond the scope of this manual to describe individual AltiVec technology implementations on PowerPC processors. It must be kept in mind that each PowerPC processor is unique in its implementation of the AltiVec technology.

The information in this book is subject to change without notice, as described in the disclaimers on the title page of this book. As with any technical documentation, it is the readers' responsibility to be sure they are using the most recent version of the documentation. For more information, contact your sales representative or visit our web site at: <http://www.mot.com/SPS/PowerPC/>.

## Audience

This manual is intended for system software and hardware developers and application programmers who want to develop products using the AltiVec technology extension to the PowerPC processors in general. It is assumed that the reader understands operating systems, microprocessor system design, and the basic principles of RISC processing.

This revision of this book describes how the AltiVec technology interacts with both the 64- and the 32-bit portions of the PowerPC architecture.

## Organization

Following is a summary and a brief description of the major sections of this manual:

- Chapter 1, "Overview," is useful for those who want a general understanding of the features and functions of the AltiVec technology. This chapter provides an overview of how the AltiVec technology defines the register set, operand conventions, addressing modes, instruction set, cache model, and exception model.

- Chapter 2, “AltiVec Register Set,” is useful for software engineers who need to understand the PowerPC programming model for the three programming environments. The chapter also discusses the functionality of the AltiVec technology registers and how they interact with the current PowerPC registers.
- Chapter 3, “Operand Conventions,” describes how the AltiVec technology interacts with the PowerPC conventions for storing data in memory, including information regarding alignment, single-precision floating-point conventions, and big- and little-endian byte ordering.
- Chapter 4, “Addressing Modes and Instruction Set Summary,” provides an overview of the AltiVec technology addressing modes and a brief description of the AltiVec technology instructions organized by function.
- Chapter 5, “Cache, Exceptions, and Memory Management,” provides a discussion of the cache and memory model defined by the VEA and aspects of the cache model that are defined by the OEA. It also describes the exception model defined in the UISA.
- Chapter 6 “Instruction Set,” functions as a handbook for the AltiVec instruction set. Instructions are sorted by mnemonic. Each instruction description includes the instruction formats and an individualized legend that provides such information as the level(s) of the PowerPC architecture in which the instruction may be found and the privilege level of the instruction. It also includes figures where it helps in understanding what the instruction does.
- Appendix A, “Instruction Set,” lists all the AltiVec instructions. Instructions are grouped according to mnemonic, opcode, and form.
- This manual also includes a glossary and an index.

## Suggested Reading

This section lists additional reading that provides background for the information in this manual as well as general information about the AltiVec technology and PowerPC architecture.

## General Information

The following documentation provides useful information about the PowerPC architecture and computer architecture in general:

- The following books are available from the Morgan-Kaufmann Publishers, 340 Pine Street, Sixth Floor, San Francisco, CA 94104; Tel. (800) 745-7323 (U.S.A.), (415) 392-2665 (International); internet address: [mkp@mkp.com](mailto:mkp@mkp.com).
  - *The PowerPC Architecture: A Specification for a New Family of RISC Processors*, Second Edition, by International Business Machines, Inc.  
 Updates to the architecture specification are accessible via the world-wide web at <http://www.austin.ibm.com/tech/ppc-chg.html>.
  - *PowerPC Microprocessor Common Hardware Reference Platform: A System*

- Architecture*, by Apple Computer, Inc., International Business Machines, Inc., and Motorola, Inc.
- *Macintosh Technology in the Common Hardware Reference Platform*, by Apple Computer, Inc.
  - *Computer Architecture: A Quantitative Approach*, Second Edition, by John L. Hennessy and David A. Patterson,
  - *Inside Macintosh: PowerPC System Software*, Addison-Wesley Publishing Company, One Jacob Way, Reading, MA, 01867; Tel. (800) 282-2732 (U.S.A.), (800) 637-0029 (Canada), (716) 871-6555 (International).
  - *PowerPC Programming for Intel Programmers*, by Kip McClanahan; IDG Books Worldwide, Inc., 919 East Hillsdale Boulevard, Suite 400, Foster City, CA, 94404; Tel. (800) 434-3422 (U.S.A.), (415) 655-3022 (International).

## PowerPC Documentation

The PowerPC documentation is organized in the following types of documents:

- User's manuals—These books provide details about individual PowerPC implementations and are intended to be used in conjunction with *The Programming Environments Manual*. *PowerPC Microprocessor Family: The Programming Environments*, Rev. 1 provides information about resources defined by the PowerPC architecture that are common to PowerPC processors. This document describes both the 64- and 32-bit portions of the architecture. MPCFPE/AD (Motorola order #)
- *Implementation Variances Relative to Rev. 1 of The Programming Environments Manual* is available via the world-wide web at <http://www.mot.com/SPS/powerpc/>.
- Addenda/errata to user's manuals—Because some processors have follow-on parts an addendum is provided that describes the additional features and changes to functionality of the follow-on part. These addenda are intended for use with the corresponding user's manuals.
- Hardware specifications—Hardware specifications provide specific data regarding bus timing, signal behavior, and AC, DC, and thermal characteristics, as well as other design considerations for each PowerPC implementation.
- Technical Summaries—Each PowerPC implementation has a technical summary that provides an overview of its features. This document is roughly the equivalent to the overview (Chapter 1) of an implementation's user's manual.
- *PowerPC Microprocessor Family: The Programmer's Reference Guide*: MPCPRG/D (Motorola order #) is a concise reference that includes the register summary, memory control model, exception vectors, and the PowerPC instruction set.
- *PowerPC Microprocessor Family: The Programmer's Pocket Reference Guide*: MPCPRGREF/D (Motorola order #): This foldout card provides an overview of the PowerPC registers, instructions, and exceptions for 32-bit implementations.

- Application notes—These short documents contain useful information about specific design issues useful to programmers and engineers working with PowerPC processors.
- Documentation for support chips

Additional literature on AltiVec technology and PowerPC implementations is being released as new processors become available. For a current list of AltiVec technology and PowerPC documentation, refer to the website at <http://www.mot.com/SPS/PowerPC/>.

## Conventions

Throughout the documentation when a register or bit is “set” it means the register or bit is set to 1, and when a register is “cleared” it means the register or bit is set to 0.

This document uses the following notational conventions:

<b>mnemonics</b>	Instruction mnemonics are shown in lowercase bold.
<i>italics</i>	Italics indicate variable command parameters, for example, <b>bcctrx</b> . Book titles in text are set in italics.
0x0	Prefix to denote hexadecimal number
0b0	Prefix to denote binary number
<b>rA, rB</b>	Instruction syntax used to identify a source GPR
<b>rD</b>	Instruction syntax used to identify a destination GPR
<b>frA, frB, frC</b>	Instruction syntax used to identify a source FPR
<b>frD</b>	Instruction syntax used to identify a destination FPR
REG[FIELD]	Abbreviations or acronyms for registers are shown in uppercase text. Specific bits, fields, or ranges appear in brackets. For example, <b>MSR[LE]</b> refers to the little-endian mode enable bit in the machine state register.
<b>vA, vB, vC</b>	Instruction syntax used to identify a source VR
<b>vD</b>	Instruction syntax used to identify a destination VR
x	In certain contexts, such as a signal encoding, this indicates a don't care.
<i>n</i>	Used to express an undefined numerical value
¬	NOT logical operator
&	AND logical operator
	OR logical operator
<b>U</b>	This symbol identifies text that is relevant with respect to the PowerPC user instruction set architecture (UISA). This symbol is used both for information that can be found in the UISA specification


as well as for explanatory information related to that programming environment.



This symbol identifies text that is relevant with respect to the PowerPC virtual environment architecture (VEA). This symbol is used both for information that can be found in the VEA specification as well as for explanatory information related to that programming environment.



This symbol identifies text that is relevant with respect to the PowerPC operating environment architecture (OEA). This symbol is used both for information that can be found in the OEA specification as well as for explanatory information related to that programming environment.



Indicates reserved bits or bit fields in a register. Although these bits may be written to as either ones or zeros, they are always read as zeros.

Additional conventions used with instruction encodings are described in Section 6.1, “Instruction Formats.”

## Acronyms and Abbreviations

Table i contains acronyms and abbreviations that are used in this document. Note that the meanings for some acronyms (such as SDR1 and XER) are historical, and the words for which an acronym stands may not be intuitively obvious.

**Table i. Acronyms and Abbreviated Terms**

Term	Meaning
ALU	Arithmetic logic unit
ASR	Address space register
BAT	Block address translation
BPU	Branch processing unit
BUID	Bus unit ID
CR	Condition register
CTR	Count register
DABR	Data address breakpoint register
DAR	Data address register
DEC	Decrementer register
DSISR	Register used for determining the source of a DSI exception
EA	Effective address
ECC	Error checking and correction

**Table i. Acronyms and Abbreviated Terms (Continued)**

<b>Term</b>	<b>Meaning</b>
FPR	Floating-point register
FPSCR	Floating-point status and control register
FPU	Floating-point unit
GPR	General-purpose register
IEEE	Institute of Electrical and Electronics Engineers
ITLB	Instruction translation lookaside buffer
IU	Integer unit
L2	Secondary cache
LIFO	Last-in-first-out
LR	Link register
LRU	Least recently used
LSB	Least-significant byte
lsb	Least-significant bit
LSQ	Least-significant quad-word
lsq	Least-significant quad-word
MESI	Modified/exclusive/shared/invalid—cache coherency protocol
MMU	Memory management unit
MSB	Most-significant byte
msb	Most-significant bit
MSQ	Most-significant quad-word
msq	Most-significant quad-word
MSR	Machine state register
NaN	Not a number
NIA	Next instruction address
No-op	No operation
OEA	Operating environment architecture
PTEG	Page table entry group
RISC	Reduced instruction set computing
RTL	Register transfer language
RWITM	Read with intent to modify
SIMM	Signed immediate value
SPR	Special-purpose register



**Table i. Acronyms and Abbreviated Terms (Continued)**

Term	Meaning
SR	Segment register
SRR0	Machine status save/restore register 0
SRR1	Machine status save/restore register 1
STE	Segment table entry
TB	Time base register
TLB	Translation lookaside buffer
UIMM	Unsigned immediate value
UISA	User instruction set architecture
VA	Virtual address
VEA	Virtual environment architecture
VR	Vector register
XATC	Extended address transfer code
XER	Register used primarily for indicating conditions such as carries and overflows for integer operations

## Terminology Conventions

Table ii lists certain terms used in this manual that differ from the architecture terminology conventions.

**Table ii. Terminology Conventions**

The Architecture Specification	This Manual
Data storage interrupt (DSI)	DSI exception
Extended mnemonics	Simplified mnemonics
Instruction storage interrupt (ISI)	ISI exception
Interrupt	Exception
Privileged mode (or privileged state)	Supervisor-level privilege
Problem mode (or problem state)	User-level privilege
Real address	Physical address
Relocation	Translation
Storage (locations)	Memory
Storage (the act of)	Access



Table iii describes instruction field notation conventions used in this manual.

**Table iii. Instruction Field Conventions**

The Architecture Specification	Equivalent to:
BA, BB, BT	<b>crbA, crbB, crbD</b> (respectively)
BF, BFA	<b>crfD, crfS</b> (respectively)
D	d
DS	ds
FLM	FM
FRA, FRB, FRC, FRT, FRS	<b>frA, frB, frC, frD, frS</b> (respectively)
FXM	CRM
RA, RB, RT, RS	<b>rA, rB, rD, rS</b> (respectively)
SI	SIMM
U	IMM
UI	UIMM
VA, VB, VT, VS	<b>vA, vB, vD, vS</b> (respectively)
VEC	AltiVec technology
/, //, ///	0...0 (shaded)

**PRELIMINARY**

# Chapter 1

## Overview

The AltiVec™ technology provides a software model that accelerates the performance of various software applications and runs on RISC (reduced instruction set computing) microprocessors. The AltiVec technology extends the instruction set architecture (ISA) of the PowerPC architecture. AltiVec technology is a short vector parallel architecture. The AltiVec ISA is based on separate vector/SIMD-style (single instruction stream, multiple data streams) execution units that have high data parallelism. That is, the AltiVec technology operations can perform on multiple data elements in a single instruction. The term ‘vector’ in this document, refers to the spatial parallel processing of short, fixed-length one-dimensional matrices performed by an execution unit. It should not be confused with the temporal parallel (pipelined) processing of long, variable-length vectors performed by classical vector machines. High degrees of parallelism are achievable with simple in-order instruction dispatch and low-instruction bandwidth. However, the ISA is designed so as not to impede additional parallelism through superscalar dispatch to multiple execution units or multithreaded execution unit pipelines.

All instructions are designed to be easily pipelined with pipeline latencies no greater than scalar, double-precision floating-point multiply-add. No instruction specifies an operation that will present a frequency limitation beyond those already imposed by existing PowerPC instructions. There are no operating mode switches which preclude fine grain interleaving of instructions with the existing floating-point and integer instructions. Parallelism with the integer and floating-point instructions is simplified by the fact that the vector unit never generates an exception and has few shared resources or communication paths that require it to be tightly synchronized with the other units. By using the SIMD parallelism, performance can be accelerated on PowerPC processors to a level that can allow concurrent real-time processing of one or more streams.

In this document the term ‘implementation’ refers to a hardware device (typically a microprocessor) that complies with the specifications defined by the architecture.

The AltiVec technology can be used as an extension to various RISC microprocessors; however, in this book it is discussed within the context of the PowerPC architecture described as follows:

- Programming model
  - Instruction set—The AltiVec instruction set specifies instructions that extend the

PowerPC instruction set. These instructions are organized similar to PowerPC instructions (such as vector load/store, vector integer, and vector floating-point instructions). The specific instructions, and the forms used for encoding them, are provided in Appendix A, “Instruction Set.”

- Register set—The AltiVec programming model defines new AltiVec registers, additions to the PowerPC register set, and how existing PowerPC registers are affected by the AltiVec technology. The model also discusses memory conventions, including details regarding the byte ordering for quad words.
- Memory model—The AltiVec technology specifies additional cache management instructions. That is, a program can execute AltiVec software instructions that indicate when a sequence of memory units (data stream/stream) are likely to be accessed.
- Exception model—To ensure efficiency, the AltiVec technology provides only an AltiVec unavailable interrupt (VUI) exception, a data storage interrupt (DSI) exception, and trace exception (if implemented). There are no exceptions other than DSI exceptions on loads and stores. The AltiVec instructions can cause PowerPC exceptions.
- Memory management model—The memory model for the AltiVec technology is the same as it is implemented for the PowerPC architecture. AltiVec memory accesses are always assumed to be aligned. If an operand is unaligned, additional AltiVec instructions are used to ensure that it is correctly placed in a vector register or in memory.
- Time-keeping model—The PowerPC time-keeping model is not impacted by the AltiVec technology.

To locate any published errata or updates for this document, refer to the website at <http://www.mot.com/SPS/PowerPC/>.

This chapter provides an overview of the major characteristics of the AltiVec technology in the order in which they are addressed in this book:

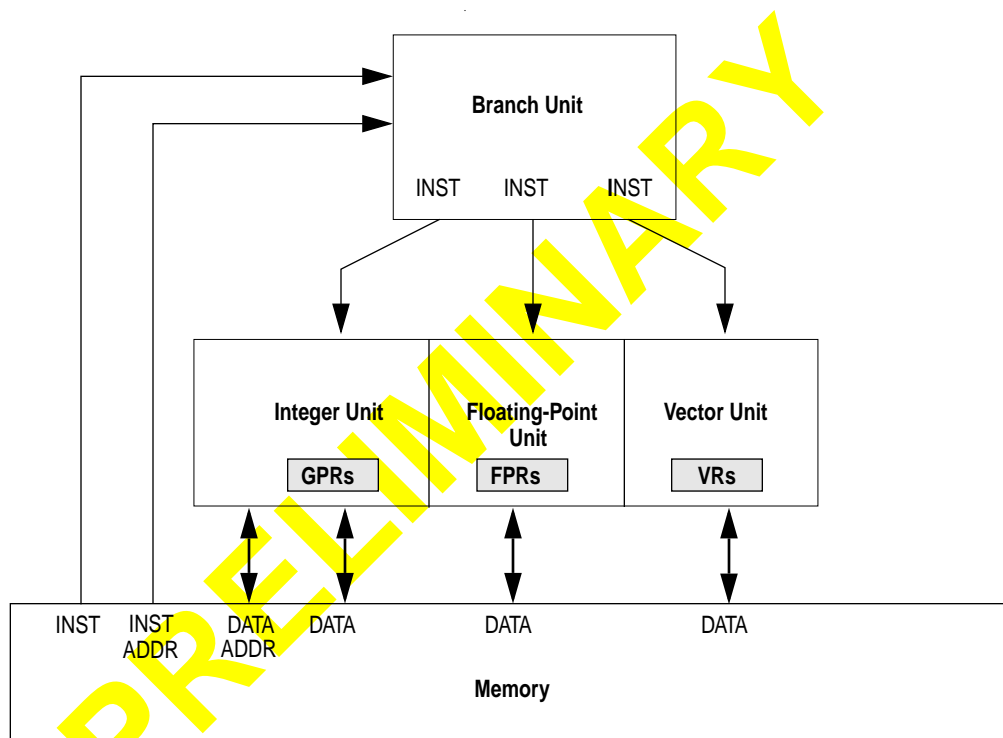
- Register set and programming model
- Instruction set and addressing modes
- Cache, exceptions, and memory management

## 1.1 Overview

The AltiVec technology’s SIMD-style extension provides an approach to accelerating the processing of data streams. Using the AltiVec instructions can provide a significant speedup for communications, multimedia, and other performance-driven applications by using data-level parallelism where available, matching scalar performance in serial sections of media applications, keeping media processing within the AltiVec unit, and minimizing bandwidth and latency memory access bottlenecks.

AltiVec technology expands the PowerPC architecture through the addition of a 128 bit vector execution unit, which operates concurrently with the existing integer- and floating-point units. A new vector execution unit provides highly parallel operations, allowing for simultaneous execution of multiple operations in a single clock cycle.

The AltiVec technology can be thought of as a set of registers and execution units that can be added to the PowerPC architecture in a manner analogous to the addition of floating point units. Floating-point units were added to provide support for high-precision scientific calculations and the AltiVec technology is added to the PowerPC architecture to accelerate the next level of performance-driven, high-bandwidth communications and computing applications. Figure 1-1 provides the high level structural overview for PowerPC with the AltiVec technology.



**Figure 1-1. High Level Structural Overview of PowerPC with AltiVec Technology**

The AltiVec technology is purposefully simple such that there are no exceptions other than DSI exceptions on loads and stores, no hardware unaligned access support, and no complex functions. The AltiVec technology is scaled down to only the necessary pieces in order to facilitate efficient cycle time, latency, and throughput on hardware implementations.

The AltiVec technology defines the following:

- Fixed 128-bit wide vector length that can be subdivided into sixteen 8-bit bytes, eight 16-bit half words, or four 32-bit words
- Vector register file (VRF) architecturally separate from floating-point registers (FPRs) and general-purpose registers (GPRs)

- Vector integer and floating-point arithmetic
- Four operands for most instructions (three source operands and one result)
- Saturation clamping, (that is, unsigned results are clamped to zero on underflow and to the maximum positive integer value ( $2^{n-1}-1$ , for example, 255 for byte fields) on overflow. For signed results, saturation clamps results to the smallest representable negative number ( $-2^{n-1}-1$ , for example, 128 for byte fields) on underflow, and to the largest representable positive number ( $2^{n-1}-1$ , for example, +127 for byte fields) on overflow)
- No mode switching that would increase the overhead of using the instructions
- Operations selected based on utility to digital signal processing algorithms (including 3D).
- AltiVec instructions provide a vector compare and select mechanism to implement conditional execution as the preferred way to control data flow in AltiVec programs
- Enhanced cache/memory interface

The AltiVec ISA supports the following:

- Audio decode and encode: G.711, G.721, G.723 and AC-3
- Communications
  - Software modem: FAX, V.32, V.34, Cable
  - Data encryption: DES, RSA
- 2D and 3D graphics: QuickDraw, OpenGL, VRML, Games, High-precision CAD
- High-fidelity audio: 3D audio, AC-3
- JPEG image processing
- Java
- Motion video decode and encode: MPEG-1, MPEG-2, MPEG-4, and H.234
- Real-time continuous speech I/O: HMM, Viterbi acceleration, Neural algorithms
- Video conferencing: H.261, H.263

### 1.1.1 The 64-Bit AltiVec Technology and the 32-Bit Subset




The AltiVec technology supports the following modes of PowerPC operations:

- 64-bit implementations/64-bit mode—The AltiVec technology defines interactions with the PowerPC 64-bit registers.
- 64-bit implementations/32-bit mode—The AltiVec technology defines interaction with the conventions for 32-bit implementations of PowerPC registers.

For further details on the 64-bit PowerPC architecture and the 32-bit subset refer to Chapter 1, “Overview,” in the *PowerPC Microprocessor Family: The Programming Environments Manual*.

### 1.1.2 The Levels of the AltiVec ISA

The AltiVec ISA follows the layering of PowerPC architecture. The PowerPC architecture has three levels defined as follows:

- PowerPC user instruction set architecture (UISA) —The UISA defines the level of the architecture to which user-level (referred to as problem state in the architecture specification) software should conform. The UISA defines the base user-level instruction set, user-level registers, data types, floating-point memory conventions and exception model as seen by user programs, and the memory and programming models. The icon shown in the margin identifies text that is relevant to the UISA. 
- PowerPC virtual environment architecture (VEA)—The VEA defines additional user-level functionality that falls outside typical user-level software requirements. The VEA describes the memory model for an environment in which multiple devices can access memory, defines aspects of the cache model, defines cache control instructions, and defines the time base facility from a user-level perspective. The icon shown in the margin identifies text that is relevant to the VEA. 
- PowerPC operating environment architecture (OEA)—The OEA defines supervisor-level (referred to as privileged state in the architecture specification) resources typically required by an operating system. The OEA defines the PowerPC memory management model, supervisor-level registers, synchronization requirements, and the exception model. The OEA also defines the time base feature from a supervisor-level perspective. The icon shown in the margin identifies text that is relevant to the OEA. 

The AltiVec technology defines instructions at the UISA and VEA levels. The distinctions between the levels is noted in the text throughout the document

### 1.1.3 Features Not Defined by the AltiVec ISA

Because flexibility is an important design goal of the AltiVec technology, there are many aspects of the microprocessor design, typically relating to the hardware implementation, that the AltiVec ISA does not define, for example, the number and the nature of execution units. The AltiVec ISA is a vector/SIMD architecture, and as such makes it easier to implement pipelining instructions and parallel execution units to maximize instruction throughput. However, the AltiVec ISA does not define the internal hardware details of implementations. For example, one processor may use a simple implementation having two vector execution units whereas another may provide a bigger, faster microprocessor design with several concurrently pipelined vector arithmetic logical units (ALUs) with separate load/store units (LSUs) and prefetch units.


## 1.2 The AltiVec Architectural Model

This section provides overviews of aspects defined by the AltiVec ISA, following the same order as the rest of this book. The topics are as follows:

- Registers and programming model
- Operand conventions
- Instruction set and addressing modes
- Cache model, exceptions, and memory management

### 1.2.1 AltiVec Registers and Programming Model

In the AltiVec technology, the ALU operates on from one to three source vectors and produces a single result/destination vector on each instruction. The ALU is a SIMD-style arithmetic unit that performs the same operation on all the data elements that comprise each vector. This scheme allows efficient code scheduling in a highly parallel processor. Load and store instructions are the only instructions that transfer data between registers and memory. The ALU and vector register file is shown in Figure 1-2.



**Figure 1-2. AltiVec Top-Level Diagram**

The ALU is a SIMD-style arithmetic unit in which an instruction performs the same operations in parallel on all the data elements that comprise each vector. Architecturally, the



vector register file (VRF) is separate from the GPRs and FPRs. The AltiVec programming model incorporates the 32 registers of the VRF, each register is 128 bits wide.


## 1.2.2 Operand Conventions

Operand conventions define how data is stored in vector registers and memory.

### 1.2.2.1 Byte Ordering

The default mapping for AltiVec ISA is PowerPC big-endian, but AltiVec ISA provides the option of operating in either big- or little-endian mode. The endian support of the PowerPC architecture does not address any data element larger than a double word; the basic memory unit for vectors is a quad word.

Big-endian byte ordering is shown in Figure 1-3.



**Figure 1-3. Big-Endian Byte Ordering for a Vector Register**

As shown in Figure 1-3, the elements in vector registers are numbered using big-endian byte ordering. For example, the high-order (or most significant) byte element is numbered 0 and the low-order (or least significant) byte element is numbered 15.

When defining high order and low order for elements in a vector register, be careful not to confuse its meaning based on the bit numbering. That is, in Figure 1-3 the high-order half word for word 0 (bits 0–15), would be half word 0 (bits 0–7), and the low-order half word for word 0 would be half word 1 (bits 8–15).

In big-endian mode, a n AltiVec quad word load instruction for which the effective address (EA) is quad-word aligned places the byte addressed by EA into byte element 0 of the target vector register. The byte addressed by EA + 1 is placed in byte element 1, and so forth. Similarly, an AltiVec quad word store instruction for which the EA is quad word-aligned places byte element 0 of the source vector register into the byte addressed by EA. Byte element 1 is placed into the byte addressed by EA + 1, and so forth.

### 1.2.2.2 Floating-Point Conventions

The AltiVec ISA basically has two modes for floating point, that is a Java-/IEEE-/C9X-compliant mode or a possibly faster non-Java/non-IEEE mode. AltiVec ISA conforms to the Java Language Specification 1 (hereafter referred to as Java), that is a subset of the default environment specified by the IEEE standard (ANSI/IEEE Standard 754-1985, IEEE Standard for Binary Floating-Point Arithmetic). For aspects of floating-point behavior that are not defined by Java but are defined by the IEEE standard, AltiVec ISA conforms to the IEEE standard. For aspects of floating-point behavior that are defined neither by Java nor by the IEEE standard but are defined by the C9X Floating-Point Proposal, WG14/N546 X3J11/96-010 (Draft 2/26/96) (hereafter referred to as C9X), AltiVec ISA conforms to C9X when in Java-compliant mode.

### 1.2.3 AltiVec Instruction Set and Addressing Modes

As with PowerPC instructions, AltiVec instructions are encoded as single-word (32-bit) instructions. Instruction formats are consistent among all instruction types, permitting decoding to be parallel with operand accesses. This fixed instruction length and consistent format simplifies instruction pipelining. AltiVec load, store, and stream prefetch instructions use secondary opcodes in primary opcode 31 (0b011111). AltiVec ALU-type instructions use primary opcode point 4 (0b000100).

AltiVec ISA supports both intraelement and interelement operations. In an intraelement operation, elements work in parallel on the corresponding elements from multiple source operand registers and place the results in the corresponding fields in the destination operand register. An example of an intraelement operation is the Vector Add Signed Word Saturate (**vaddsws**) instruction shown in Figure 1-4.





Figure 1-4. Intraelement Example, **vaddsws**

In this example, the four signed integer (32 bits) elements in register **vA** are added to the corresponding four signed integer (32 bits) elements in register **vB** and the four results are placed in the corresponding elements in register **vD**.

In interelement operations data paths cross over. That is, different elements from each source operand are used in the resulting destination operand. An example of an interelement operation is the Vector Permute (**vperm**) instruction shown in Figure 1-5.



**Figure 1-5. Interelement Example. vperm**

In this example, **vperm** allows any byte in two source vector registers (**vA** and **vB**) to be copied to any byte in the destination vector register, **vD**. The bytes in a third source vector register (**vC**) specify from which byte in the first two source vector registers the corresponding target byte is to be copied. In this case the elements from the source vector registers do not have corresponding elements that operate on the destination register.

Most arithmetic and logical instructions are intraelement operations. The data paths for the ALU run primarily north and south with little crossover. The crossover data paths have been restricted as much as possible to the interelement manipulation instructions (unpack, pack, permute, etc.) with a vision toward implementing the ALU and shift/permute networks as separate execution units. The following list of instructions distinguishes between interelement and intraelement instructions:

- Vector intraelement instructions
  - Vector integer instructions
    - Vector integer arithmetic instructions
    - Vector integer compare instructions
    - Vector integer rotate and shift instructions
  - Vector floating-point instructions
    - Vector floating-point arithmetic instructions
    - Vector floating-point rounding and conversion instructions
    - Vector floating-point compare instruction
    - Vector floating-point estimate instructions
  - Vector memory access instructions
- Vector interelement instructions
  - Vector alignment support instructions
  - Vector permutation and formatting instructions
    - Vector pack instructions
    - Vector unpack instructions

- Vector merge instructions
- Vector splat instructions
- Vector permute instructions
- Vector shift left/right instructions

### 1.2.3.1 Altivec Instruction Set

Although these categories are not defined by the Altivec ISA, the Altivec instructions can be grouped as follows:



- Vector integer arithmetic instructions—These instructions are defined by the UISA. They include computational, logical, rotate, and shift instructions.
  - Vector integer arithmetic instructions
  - Vector integer compare instructions
  - Vector integer logical instructions
  - Vector integer rotate and shift instructions
- Vector floating-point arithmetic instructions—These include floating-point arithmetic instructions defined by the UISA.
  - Vector floating-point arithmetic instructions
  - Vector floating-point multiply/add instructions
  - Vector floating-point rounding and conversion instructions
  - Vector floating-point compare instruction
  - Vector floating-point estimate instructions
- Vector load and store instructions—These include load and store instructions for vector registers defined by the UISA.
- Vector permutation and formatting instructions—These instructions are defined by the UISA.
  - Vector pack instructions
  - Vector unpack instructions
  - Vector merge instructions
  - Vector splat instructions
  - Vector permute instructions
  - Vector select instructions
  - Vector shift instructions
- Processor control instructions—These instructions are used to read and write from the Altivec status and control register (VSCR). These instructions are defined by the UISA.
- Memory control instructions—These instructions are used for managing of caches (user level and supervisor level). The instructions are defined by VEA.



### 1.2.4 AltiVec Cache Model

U

The AltiVec ISA defines several instructions for enhancements to cache management. These instructions allow software to indicate to the cache hardware how it should prefetch and prioritize writeback of data. The AltiVec ISA does not define hardware aspects of cache implementations.

V

### 1.2.5 AltiVec Exception Model

The AltiVec vector unit never generates an exception. Data stream instructions will never cause an exception themselves. Therefore, on any event that would cause an exception on a normal load or store, such as a page fault or protection violation, the data stream instruction does not take a data storage interrupt (DSI) exception; instead, it simply aborts and is ignored. Most AltiVec instructions do not generate any kind of exception. Vector load and store instructions that attempt to access a direct-store segment will cause a DSI exception.

The AltiVec unit does not report IEEE exceptions; there are no status flags and the unit has no architecturally visible traps. Default results are produced for all exception conditions as specified first by the Java specification. If no default exists, the IEEE standard's default is used. Then, if no default exists, the C9X default is used.

### 1.2.6 Memory Management Model

In a PowerPC processor the MMU's primary functions are to translate logical (effective) addresses to physical addresses for memory accesses and I/O accesses (most I/O accesses are assumed to be memory-mapped) and to provide access protection on a block or page basis. Some protection is also available even if translation is disabled. Typically, it is not programmable. The AltiVec ISA does not provide any additional instructions to the PowerPC memory management model, but the AltiVec instructions have options to ensure that an operand is correctly placed in a vector register or in memory.

**PRELIMINARY**

# Chapter 2


## Altivec Register Set

This chapter describes the register organization defined by the Altivec technology. It also describes how Altivec instructions affect some of the PowerPC registers. The Altivec ISA defines register-to-register operations for all computational instructions. Source data for these instructions is accessed from the on-chip vector registers (VRs) or are provided as immediate values embedded in the opcode. Architecturally, the VRs are separate from the general-purpose registers (GPRs) and floating-point registers (FPRs). Data is transferred between memory and vector registers with explicit Altivec load and store instructions only.

Note that the handling of reserved bits in any register is implementation-dependent. Software is permitted to write any value to a reserved bit in a register. However, a subsequent reading of the reserved bit returns 0 if the value last written to the bit was 0 and returns an undefined value (may be 0 or 1) otherwise. This means that even if the last value written to a reserved bit was 1, reading that bit may return 0.

### 2.1 Altivec Register File

The Altivec vector register file (VRF), shown in Figure 2-1, has 32 registers, each is 128 bits wide. Each vector register can hold sixteen 8-bit elements, eight 16-bit elements, or four 32-bit elements.




**Figure 2-1. AltiVec Register File**

The vector registers are accessed as vector instruction operands. Access to registers are explicit as part of the execution of an instruction.

### 2.1.1 The Vector Status and Control Register (VSCR)

The vector status and control register (VSCR) is a special 32-bit vector register (not an SPR) that is read and written in a manner similar to the FPSCR in the PowerPC scalar floating-point unit. The VSCR is shown in Figure 2-2.





**Figure 2-2. Vector Status and Control Register (VSCR)**

The VSCR has two defined bits, the AltiVec non-Java mode (NJ) bit (VSCR[15]) and the AltiVec saturation (SAT) bit (VSCR[31]); the remaining bits are reserved.

Special instructions Move from Vector Status and Control Register (**mfvsr**) and Move to Vector Status and Control Register (**mtvsr**) are provided to move the VSCR from and to a vector register. When moved to or from a vector register, the 32-bit VSCR is right-justified in the 128-bit vector register. When moved to a vector register, the upper 96 bits VRx[0–95] of the vector register are cleared, so the VSCR in a vector register looks as shown in Figure 2-3.



 Reserved



**Figure 2-3. VSCR Moved to a Vector Register**

VSCR bit settings are shown in Table 2-1.

**Table 2-1. VSCR Field Descriptions**

Bit(s)	Name	Description
0–14	—	Reserved. The handling of reserved bits is the same as the normal PowerPC implementation, that is, system registers such as XER and FPSCR are implementation-dependent. Software is permitted to write any value to such a bit. A subsequent reading of the bit returns 0 if the value last written to the bit was 0 and returns an undefined value (0 or 1) otherwise.
15	NJ	<p>Non-Java. A mode control bit that determines whether AltiVec floating-point operations will be performed in a Java-IEEE-C9X-compliant mode or a possibly faster non-Java/non-IEEE mode.</p> <p>NJ = 0 = off, the Java-IEEE-C9X-compliant mode is selected. Denormalized values are handled as specified by Java, IEEE, and C9X standard.</p> <p>NJ = 1 = on, the non-Java/non-IEEE-compliant mode is selected. If an element in a source vector register contains a denormalized value, the value 0 is used instead. If an instruction causes an underflow exception, the corresponding element in the target VR is cleared to 0. In both cases the 0 has the same sign as the denormalized or underflowing value.</p> <p>This mode is described in detail in the floating-point overview Section 3.2.1, “Floating-Point Modes.”</p>
16–30	—	Reserved. The handling of reserved bits is the same as the normal PowerPC implementation, that is, system registers such as XER and FPSCR are implementation-dependent. Software is permitted to write any value to such a bit. A subsequent reading of the bit returns 0 if the value last written to the bit was 0 and returns an undefined value (0 or 1) otherwise.

**Table 2-1. VSCR Field Descriptions (Continued)**

Bit(s)	Name	Description
31	SAT	<p>Saturation. A sticky status bit indicating that some field in a saturating instruction saturated since the last time SAT was cleared. In other words when SAT = 1 it remains set to 1 until it is cleared to 0 by an <b>mtvscr</b> instruction. For further discussion refer to Section 4.2.1.1, “Saturation Detection.”</p> <p>1 = The AltiVec saturate instruction implicitly sets when saturation has occurred on the results one of the AltiVec instructions having saturate in its name:  Move To VSCR (<b>mtvscr</b>)  Vector Add Integer with Saturation (<b>vaddubs</b>, <b>vadduhs</b>, <b>vadduws</b>, <b>vaddubs</b>, <b>vaddshs</b>, <b>vaddsws</b>)  Vector Subtract Integer with Saturation (<b>vsububs</b>, <b>vsubuhs</b>, <b>vsubuws</b>, <b>vsububs</b>, <b>vsubshs</b>, <b>vsubsws</b>)  Vector Multiply-Add Integer with Saturation (<b>vmhaddshs</b>, <b>vmhraddshs</b>)  Vector Multiply-Sum with Saturation (<b>vmsumuhs</b>, <b>vmsumshs</b>, <b>vmsumsws</b>)  Vector Sum-Across with Saturation (<b>vsumsws</b>, <b>vsum2sws</b>, <b>vsum4sbs</b>, <b>vsum4shs</b>, <b>vsum4ubs</b>)  Vector Pack with Saturation (<b>vpkuhus</b>, <b>vpkuwus</b>, <b>vpkshus</b>, <b>vpkswus</b>, <b>vpkshss</b>, <b>vpkswss</b>)  Vector Convert to Fixed-Point with Saturation (<b>vctuxs</b>, <b>vctxs</b>)</p> <p>0 = Indicates no saturation occurred, <b>mtvscr</b> can explicitly clear this bit.</p>

The **mtvscr** is context synchronizing. This implies that all AltiVec instructions logically preceding an **mtvscr** in the program flow will execute in the architectural context (NJ mode) that existed prior to completion of the **mtvscr**, and that all instructions logically following the **mtvscr** will execute in the new context (NJ mode) established by the **mtvscr**.

After an **mtvscr** instruction executes, the result in the target vector register will be architecturally precise. That is, it will reflect all updates to the SAT bit that could have been made by vector instructions logically preceding it in the program flow, and further, it will not reflect any SAT updates that may be made to it by vector instructions logically following it in the program flow. Reading the VSCR can be much slower than typical AltiVec instructions, and therefore care must be taken in reading it to avoid performance problems.

### 2.1.2 VRSAGE Register (VRSAGE)

The VRSAGE vector register is a separate register used to assist in application and operating system software in saving and restoring the architectural state across process context-switched events. VRSAGE is a new user-mode accessible 32-bit special-purpose register (SPR 256) that is added to the PowerPC architecture to assist software in providing efficient save and restore operations. The VRSAGE register (VRSAGE) is entirely maintained and managed by software, VRSAGE is shown in Figure 2-4.

VR0	VR1	VR2	VR3	VR4	VR5	VR6	VR7	VR8	VR9	VR10	VR11	VR12	VR13	VR14	VR15	VR16	VR17	VR18	VR19	VR20	VR21	VR22	VR23	VR24	VR25	VR26	VR27	VR28	VR29	VR30	VR31
0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31

**Figure 2-4. Saving/Restoring the AltiVec Context Register (VRSAGE)**

VRSAVE bit settings are shown in Table 2-2.


**Table 2-2. VRSAVE Bit Settings**

Bit(s)	Name	Description
0-31	VR $n$	1 = VR $n$ is live, it is using VR0 of the VRF for the current running process 0 = VR $n$ is dead, it is not being used for the current running process

The VRSAVE register is read or written only as the direct result of a **mfspr** or **mtspr** instruction, respectively. The recommended usage of VRSAVE is for each bit in this register to correspond to one of the vector registers and its values indicate whether the corresponding register is currently live (1) or dead (0). A live register contains data that is currently in use by the executing process, a dead register does not contain data. If VRSAVE is used to indicate which vector registers (VRs) are being used by a program, the operating system could save only those VRs when an exception occurs, and could restore only those VRs when resuming from the exception. If this approach is taken it must be applied rigorously; if a program fails to indicate that a given VR is in use, software errors may occur that will be difficult to detect and correct because they are timing-dependent. Some operating systems save and restore VRSAVE only for programs that also use other AltiVec registers.

### 2.1.3 PowerPC Condition Register

The PowerPC condition register (CR) is a 32-bit register that reflects the result of certain operations and provides a mechanism for testing and branching. For AltiVec ISA, the CR6 field can optionally be used, that is if an AltiVec instruction field's record bit (Rc) is set in a vector compare instruction, the CR6 field is updated. The bits in the PowerPC CR are grouped into eight 4-bit fields, CR0–CR7, as shown in Figure 2-5.



**Figure 2-5. Condition Register (CR)**

For more details on the CR see Chapter 2 “PowerPC Register Set,” in *PowerPC: The Programming Environments Manual*.

To control program flow based on vector data, all vector compare instructions can optionally update CR6. If the instruction field's record bit (Rc) is set in a vector compare instruction, the CR6 field is updated according to Table 2-3.

CR Bit	CR6 Field Bit	Vector Compare	Vector Compare Bounds
24	0	1 = Relation is true for all element pairs	0
25	1	0	0
26	2	1 = Relation is false for all element pairs 0 = All fields were in bounds	1 = All fields are in bounds for the <b>vcmpbfp</b> instruction so the result code of all fields is 0b00 0 = One of the fields is out of bounds for the <b>vcmpbfp</b> instruction
27	3	0	0

The Rc bit should be used sparingly. As for other PowerPC instructions, in some implementations instructions with Rc bit = 1 could have somewhat longer latency or be more disruptive to instruction pipeline flow than instructions with Rc bit = 0. Therefore techniques of accumulating results and testing infrequently are advised.

An “AltiVec Available” bit is added to the PowerPC machine state register (MSR). The VEC bit 39 is added to the PowerPC machine state register (MSR) as shown in Figure 2-6.



### Figure 2-7. Machine State Register (MSR)—32-Bit Implementation

In 32-bit PowerPC implementations, bit 6, VEC, is added to the MSR as shown in Figure 2-7. Also AltiVec data stream prefetching instructions will be suspended and resumed based on MSR[PR] and MSR[DR]. The Data Stream Touch (**dst**) and Data Stream Touch for Store (**dstst**) instructions are supported whenever MSR[DR] = 1. If either instruction is executed when MSR[DR] = 0 (real addressing mode), the results are boundedly undefined. For each existing data stream, prefetching is enabled if the MSR[DR] = 1 and MSR[PR] bit has the value it had when the **dst** or **dstst** instruction that specified the data stream was executed. Otherwise prefetching for the data stream is suspended. In particular, the occurrence of an exception suspends all data stream prefetching.

Table 2-4 shows the AltiVec bit definitions for the MSR as well as how the PR and DR bits are affected by the AltiVec data stream instructions.

**Table 2-4. MSR Bit Settings Affected by AltiVec**

Bit(s)		Name	Description
64 Bit	32 Bit		
38	6	VEC	<p>AltiVec Available</p> <p>0 When the bit is cleared to zero, the processor executes an “AltiVec Unavailable Exception” when any attempt to execute a vector instruction that accesses the vector register file (VRF) or VSCR register.</p> <p>1 The VRF and VSCR registers are accessible to vector instructions.</p> <p>Note: the VRSAVE register is not protected by MSR [VEC]. The data streaming family of instructions (<b>dst</b>, <b>dstt</b>, <b>dstst</b>, <b>dststt</b>, <b>dss</b>, and <b>dssall</b>) are not affected by the MSR[VEC], that is, the VRF and VSCR registers are available to the data streaming instructions even when the MSR[VEC] is cleared.</p>
49	17	PR	<p>Privilege level</p> <p>0 The processor can execute both user- and supervisor-level instructions.</p> <p>1 The processor can only execute user-level instructions.</p> <p>Note: Care should be taken if data-stream prefetching is used in privileged state (MSR[PR] = 0). For each existing data stream, prefetching is enabled if (a) MSR[DR] = 1 and (b) MSR[PR] has the value it had when the <b>dst</b> or <b>dstst</b> instruction that specified the data stream was executed. Otherwise, prefetching for the data stream is suspended.</p>
59	27	DR	<p>Data address translation</p> <p>0 Data address translation is disabled. If data stream touch (<b>dst</b>) and data stream touch for store (<b>dstst</b>) instructions are executed whenever DR = 0, the results are boundedly undefined.</p> <p>1 Data address translation is enabled. Data stream touch (<b>dst</b>) and data stream touch for store (<b>dstst</b>) instructions are supported whenever DR = 1.</p>

For more detailed information including the other bit settings for MSR, refer to Chapter 2, “PowerPC Register Set,” in *PowerPC Microprocessor Family: The Programming Environments Manual*.

## 2.1.5 Machine Status Save/Restore Registers (SRR)

- The machine status save/restore (SRR) registers are part of the PowerPC OEA supervisor-level registers. The SRR0 and SRR1 registers are used to save machine status on exceptions and to restore machine status when an **rfid** (or **rfi**) instruction is executed. For more detailed information, refer to Chapter 2, “PowerPC Register Set,” in *PowerPC: The Programming Environments Manual*.

### 2.1.5.1 Machine Status Save/Restore Register 0 (SRR0)

The SRR0 is a 64-bit register in 64-bit implementations and a 32-bit register in 32-bit implementations. SRR0 is used to save machine status on exceptions and restore machine status when an **rfid** (or **rfi**) instruction is executed. For the AltiVec ISA, it holds the effective address (EA) for the instruction that caused the AltiVec unavailable exception. The AltiVec unavailable exception occurs when no higher priority exception exists, and an attempt is made to execute an AltiVec instruction when MSR[VEC] = 0. The format of SRR0 is shown in Figure 2-8.

For 32-bit implementations, the format of SRR0 is that of the low-order bits (32–63) of Figure 2-8.



Figure 2-8. Machine Status Save/Restore Register 0 (SRR0)

### 2.1.6 Machine Status Save/Restore Register 1 (SRR1)

The SRR1 is a 64-bit register in 64-bit implementations and a 32-bit register in 32-bit implementations. SRR1 is used to save machine status on exceptions and to restore machine status when an **rfid** (or **rfi**) instruction is executed. The format of SRR1 is shown in Figure 2-9.




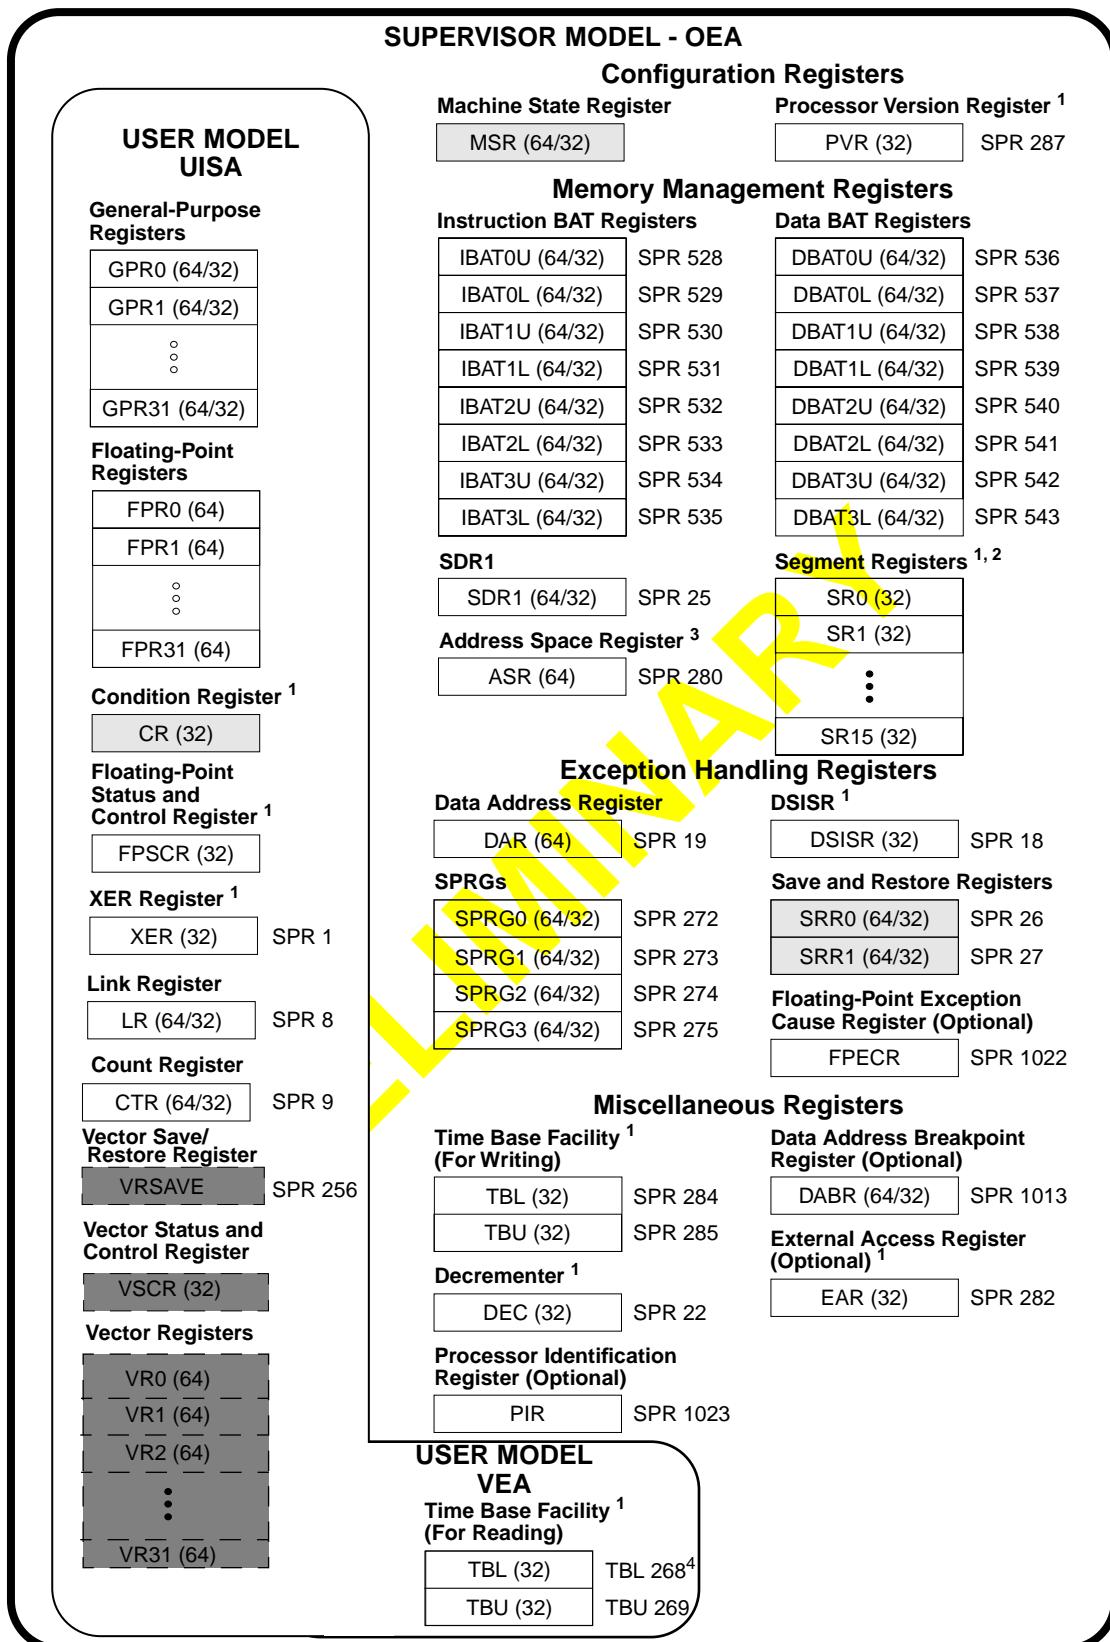
Figure 2-9. Machine Status Save/Restore Register 1 (SRR1)

In 64-bit implementations, when an AltiVec unavailable exception occurs, SRR1[33–36] and SRR1[42–47] are cleared to zero and bits MSR[0], MSR[48–55], MSR[57–59], and MSR[62–63] are placed into the corresponding bit positions of SRR1 as it was just prior to the exception. For 32-bit implementations, when an AltiVec unavailable exception occurs, SRR1[1–4] and SRR1[10–15] are cleared to zero and MSR[16–23], MSR[25–27], and MSR[30–31] are placed into the corresponding bit positions of SRR1 as they were before the exception.

## 2.2 PowerPC Register Set

The addition of the AltiVec technology adds some additional new registers as well as affecting bit settings in some of the PowerPC registers when AltiVec instructions are executed. Figure 2-10 shows a graphic representation of the entire PowerPC register set and how the AltiVec register set resides within the PowerPC architecture. The PowerPC registers affected by AltiVec instructions are shaded.

**PRELIMINARY**



<sup>1</sup> These registers are 32-bit registers only.

<sup>2</sup> These registers are on 32-bit implementations only.

<sup>3</sup> These registers are on 64-bit implementations only.

<sup>4</sup> In 64-bit implementations, TBR268 is read as a 64-bit value.

**Figure 2-10. OEA Programming Model—All Registers**



# Chapter 3

## Operand Conventions

This section describes the operand conventions as they are represented in the AltiVec technology at the UISA level. Detailed descriptions are provided of conventions used for transferring data between vector registers and memory and representing data in these vector registers using both big- and little-endian byte ordering. Additionally, the floating-point default conditions for exceptions are described.

### 3.1 Data Organization in Memory and Data Transfers

AltiVec instruction set architecture (ISA) follows the same data organization as the PowerPC architecture UISA with a few extensions. In addition to supporting byte, half-word and word operands, as defined in the PowerPC architecture UISA, AltiVec ISA supports quad-word (128-bit) operands.

The following sections describe the concepts of alignment and byte ordering of data for quad words, otherwise alignment is the same as described in Chapter 3, “Operand Conventions,” in the *PowerPC Microprocessor Family: The Programming Environments Manual*.

#### 3.1.1 Aligned and Misaligned Accesses

Vectors are accessed from memory with instructions such as Vector Load Indexed (**lvx**) and Store Vector Indexed (**stvx**) instructions. The operand of a vector register to memory access instruction has a natural alignment boundary equal to the operand length. In other words, the natural address of an operand is an integral multiple of the operand length. A memory operand is said to be aligned if it is aligned at its natural boundary; otherwise it is misaligned. AltiVec instructions are four bytes long and word-aligned as it is for PowerPC instructions.

Operands for vector register to memory access instructions have the characteristics shown in Table 3-1.

**Table 3-1. Memory Operand Alignment**

Operand	Length	Aligned Address (28-31)
Byte	8 bits (1 byte)	xxxx
Half word	2 bytes	xxx0
Word	4 bytes	xx00
Quad word	16 bytes	0000

**Note:** An x in an address bit position indicates that the bit can be 0 or 1 independent of the state of other bits in the address.

The concept of alignment is also applied more generally to data in memory. For example, an 8-byte data item is said to be half-word-aligned if its address is a multiple of two; that is, the effective address (EA) points to the next effective address that is 2 bytes (a half word) past the current effective address, that would be the  $EA + 2$  bytes, and then the next being the  $EA + 4$  bytes, and effective address would continue skipping every 2 bytes (2 bytes = 1 half word). This ensures that the effective address is half-word aligned as it points to each successive half word in memory.

It is important to understand that AltiVec memory operands are assumed to be aligned, and AltiVec memory accesses are performed as if the appropriate number of low-order bits of the specified effective address were zero. This assumption is different from PowerPC integer and floating-point memory access instructions where alignment is not always assumed. So for AltiVec ISA, the low-order bit of the effective address is ignored for half-word AltiVec memory access instructions, and the low-order four bits of the effective address are ignored for quad-word AltiVec memory access instructions. The effect is to load or store the memory operand of the specified length that contains the byte addressed by the effective address.

If a memory operand is misaligned, additional instructions must be used to correctly place the operand in a vector register or in memory. AltiVec technology provides instructions to shift and merge the contents of two vector registers. These instructions facilitate copying misaligned quad-word operands between memory and the vector registers.

### 3.1.2 AltiVec Byte Ordering

For PowerPC and AltiVec processors, the smallest addressable memory unit is the byte (8 bits), and scalars are composed of one or more sequential bytes. The AltiVec ISA supports both big- and little-endian byte ordering. The default byte ordering is big-endian. However, the code sequence used to switch from big- to little-endian mode may differ among processors.

The PowerPC architecture uses the MSR for specifying byte ordering—LE (little-endian mode). The MSR[LE] specifies the endian mode in which the processor is currently operating. A value of 0 specifies big-endian mode and a value of 1 specifies little-endian mode. For further details on PowerPC byte ordering, refer to Chapter 3, “Operand

Conventions,” in the *PowerPC Microprocessor Family: The Programming Environments Manual*

AltiVec ISA follows the endian support of PowerPC for elements up to double words. AltiVec ISA also supports quad words and additional support is provided for this. In AltiVec ISA when a 64-bit scalar is moved from a register to memory, it occupies eight consecutive bytes in memory and a decision must be made regarding byte ordering in these eight addresses.

The default byte ordering for AltiVec ISA is big-endian.

### 3.1.2.1 Big-Endian Byte Ordering

For big-endian scalars, the most-significant byte (MSB) is stored at the lowest (or starting) address while the least-significant byte (LSB) is stored at the highest (or ending) address. This is called big-endian because the big end of the scalar comes first in memory.

### 3.1.2.2 Little-Endian Byte Ordering

For little-endian scalars, the least-significant byte (LSB) is stored at the lowest (or starting) address while the MSB is stored at the highest (or ending) address. This is called little-endian because the little end of the scalar comes first in memory.

### 3.1.3 Quad Word Byte Ordering Example

The idea of big- and little-endian byte ordering is best illustrated in an example of a quad word such as 0x2021\_2223\_2425\_2627\_2829\_2A2B\_2C2D\_2E2F located in memory. This quad word is used throughout this section to demonstrate how the bytes that comprise a quad word are mapped into memory.

The quad word (0x2021\_2223\_2425\_2627\_2829\_2A2B\_2C2D\_2E2F) is shown in big-endian mapping in Figure 3-1. A hexadecimal representation is used for showing address values and the values in the contents of each byte. The address is shown below each byte's contents. The big-endian model addresses the quad word at address 0x00, which is the MSB (0x20), proceeding to the address 0x0F, which contains the LSB (0x2F).

Byte	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
	Quad Word															
Contents	20	21	22	23	24	25	26	27	28	29	2A	2B	2C	2D	2E	2F
Address	00	01	02	03	04	05	06	07	08	09	0A	0B	0C	0D	0E	0F
	↑															↑
	MSB							LSB								

**Figure 3-1. Big-Endian Mapping of a Quad Word**

Figure 3-2 shows the same quad word using little-endian mapping. In the little-endian model, the quad word's 0x00 address specifies the LSB (0x2F) and proceeds to address 0x2F which contains its MSB (0x20).

Byte	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
	Quad Word															
Contents	2F	2E	2D	2C	2B	2A	29	28	27	26	25	24	23	22	21	20
Address	00	01	02	03	04	05	06	07	08	09	0A	0B	0C	0D	0E	0F
	↑															↑
	LSB							MSB								

**Figure 3-2. Little-Endian Mapping of a Quad Word**

Figure 3-2 shows the sequence of bytes laid out with addresses increasing from left to right. Programmers familiar with little-endian byte ordering may be more accustomed to viewing quad words laid out with addresses increasing from right to left, as shown in Figure 3-3.

Byte	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
	Quad Word															
Contents	20	21	22	23	24	25	26	27	28	29	2A	2B	2C	2D	2E	2F
Address	0F	0E	0D	0C	0B	0A	09	08	07	06	05	04	03	02	01	00
	↑															↑
	MSB							LSB								

**Figure 3-3. Little-Endian Mapping of Quad Word—Alternate View**

This allows the little-endian programmer to view each scalar in its natural byte order of MSB to LSB. However, to demonstrate how the AltiVec ISA provides both big- and little-endian support, this section uses the convention of showing addresses increasing from left to right, as in Figure 3-2.

### 3.1.4 Aligned Scalars in Little-Endian Mode

The effective address (EA) calculation for the load and store instructions is described in Chapter 4, “Addressing Modes and Instruction Set Summary.” For PowerPC processors in little-endian mode, the effective address is modified before being used to access memory. In PowerPC, the three low-order address bits of the effective address are exclusive-ORed (XOR) with a three-bit value that depends on the length of the operand (1, 2, 4, or 8 bytes), as shown in Table 3-2. This address modification is called munging.

**Table 3-2 Effective Address Modifications**

Data Width (Bytes)	EA Modification
1	XOR with 0b111
2	XOR with 0b110
4	XOR with 0b100
8	No change

The munged physical address is passed to the cache or to main memory, and the specified width of the data is transferred (in big-endian order—that is, MSB at the lowest address, LSB at the highest address) between a GPR or FPR and the addressed memory locations (as modified).

Munging makes it appear to the processor that individual aligned scalars are stored as little-endian, when in fact they are stored in big-endian order but at different byte addresses within double words. Only the address is modified, not the byte order. For further details on how to align scalars in little-endian mode see Chapter 3, “Operand Conventions,” in *PowerPC: The Programming Environments Manual*.

The PowerPC address munging is performed on double-word units. In the PowerPC architecture, little-endian mode would have the double words of a quad word appear swapped. Figure 3-4 shows how a quad word appears to the processor as it munges the address when accessing memory.

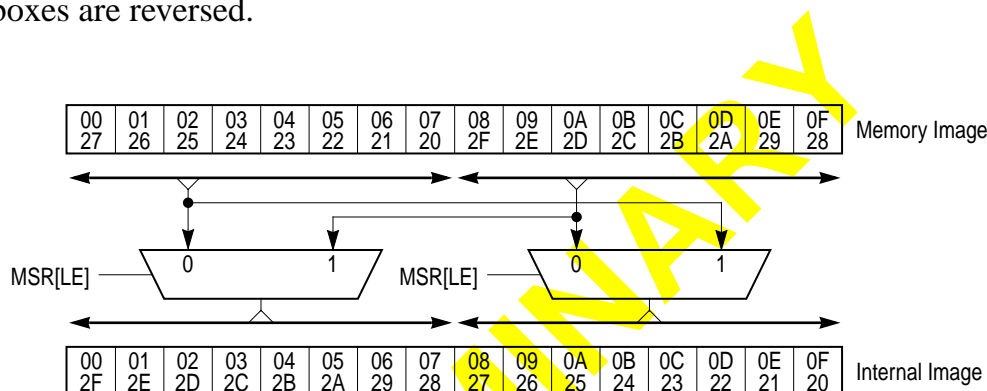
Byte	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
	Quad Word															
Contents	27	26	25	24	23	22	21	20	2F	2E	2D	2C	2B	2A	29	28
Address	00	01	02	03	04	05	06	07	08	09	0A	0B	0C	0D	0E	0F
	↑ MSB								↑ LSB							

**Figure 3-4. Quad Word Load with Munged Little-Endian Applied**

Note that double words are swapped. The byte element addressed by the quad word's base address, 0x0F, contains 0x28, while its MSB at address 0x0 contains 0x27. This is due to the PowerPC munging being applied to offsets within double words; AltiVec ISA requires a munge within quad words.

To accommodate the quad-word operands, the PowerPC architecture can not simply be extended by munging an extra address bit. It would break existing code and/or platforms. Processors that implement AltiVec technology could not be mixed with non-AltiVec processors. Instead, AltiVec processors implement a double-word swap when moving quad words between vector registers and memory.

Figure 3-5 shows how this swapping could be implemented. This diagram represents the load path double-word swapping; the store path looks the same, except that the memory and internal boxes are reversed.



**Figure 3-5. AltiVec Little Endian Swap**

In the diagram, the numbers at the top of the byte boxes represent the offset address of that byte; the numbers at the bottom are the values of the bytes at that offset. The little-endian ordering is discontinuous because the PowerPC munging is performed only on double-word units. The purpose of the double word swap within the AltiVec unit is to perform an additional swap that is not part of the PowerPC architecture.

When MSR[LE] = 1, double words are swapped and the bytes now appear in their expected ordering. When MSR[LE] = 0, no swapping is done.

To summarize, in little-endian mode, the load vector element indexed instructions (**lvebx**, **lvehx**, **lviewx**) and the store vector element indexed instructions (**stvebx**, **stvehx**, **stviewx**) have the same 3-bit address munge applied to the memory address as is specified by the PowerPC architecture for integer and floating-point loads and stores. For the quad word load vector indexed instructions (**lvx**, **lvxl**) and the store vector indexed instructions (**stvx**, **stvxl**) the two double words of the quad-word scalar data are munged and swapped as they are moved between the vector register and memory.

### 3.1.5 Vector Register and Memory Access Alignment

When loading an aligned byte, half word, or word memory operand into a vector register, the element that receives the data is the element that would have received the data had the entire aligned quad word containing the memory operand addressed by the effective address been loaded. Similarly, when an element in a vector register is stored into an aligned memory operand, the element selected to be stored is the element that would have been stored into the memory operand addressed by the effective address had the entire vector register been stored to the aligned quad word containing the memory operand addressed by the effective address. The position of the element in the target or source vector register depends on the endian mode, as described above. (Byte memory operands are always aligned.)

For aligned byte, half word, and word memory operands, if the corresponding element number is known when the program is written, the appropriate vector splat and vector permute instructions can be used to copy or replicate the data contained in the memory operand after loading the operand into a vector register. An example of this is given in detail in Section 3.1.6, “Quad-Word Data Alignment.” Another example is to replicate the element across an entire vector register before storing it into an arbitrary aligned memory operand of the same length; the replication ensures that the correct data is stored regardless of the offset of the memory operand in its aligned quad word in memory.

Since vector loads and stores are size-aligned, application binary interfaces (ABIs) should specify, and programmers should take care to align data on even quad-word boundaries for maximum performance.

### 3.1.6 Quad-Word Data Alignment

The AltiVec ISA does not provide for alignment exceptions for loading and storing data. When performing vector loads and stores, the effect is as if the low-order four bits of the address are 0x0, regardless of the actual effective address generated. Since vectors may often be misaligned due to the nature of the algorithm, AltiVec ISA provides support for post-alignment of quad-word loads and pre-alignment for quad-word stores. Note that in the following diagrams, the effect of the swapping described above is assumed and the memory diagrams will be with respect to the logical mapping of the data.

Figure 3-6 and Figure 3-7 show misaligned vectors in memory for both big- and little-endian ordering. The big-endian example assumes that the desired vector begins at address 0x03. The little-endian example assumes the desired vector begins at address 0x0D.



Byte	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
	Quad Word HI															
Contents				20	21	22	23	24	25	26	27	28	29	2A	2B	2C
Address	00	01	02	03	04	05	06	07	08	09	0A	0B	0C	0D	0E	0F
	Quad Word LO															
Contents	2D	2E	2F													
Address	10	11	12	13	14	15	16	17	18	19	1A	1B	1C	1D	1E	1F
	↑															↑
	MSB															LSB

**Figure 3-6. Misaligned Vector in Big-Endian Addressing Mode**

Figure 3-7 shows an misaligned vector in little-endian mode.

Byte	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
	Quad Word HI															
Contents														2F	2E	2D
Address	00	01	02	03	04	05	06	07	08	09	0A	0B	0C	0D	0E	0F
	Quad Word LO															
Contents	2C	2B	2A	29	28	27	26	25	24	23	22	21	20			
Address	10	11	12	13	14	15	16	17	18	19	1A	1B	1C	1D	1E	1F
	↑															↑
	LSB															MSB


**Figure 3-7. Misaligned Vector in Little-Endian Addressing Mode**

Figure 3-6 and Figure 3-7 show how such misaligned data causes data to be split across aligned quad words; only aligned quad words are loaded and/or stored by AltiVec load/store instructions. To align this vector, a program must load both (aligned) quad words that contain a portion of the misaligned vector data and then execute a Vector Permute (**vperm**) instruction to align the result.

### 3.1.6.1 Accessing a Misaligned Quad Word in Big-Endian Mode

Figure 3-1 shows the big-endian alignment model, using the example in Figure 3-6. In Figure 3-8, **vHI** and **vLO** (HI = high order, LO = low order) represent vector registers that contain the misaligned quad words containing the MSBs and LSBs, respectively, of the misaligned quad word; **vD** is the target vector register.





**Figure 3-8. Big-Endian Quad Word Alignment**

Alignment is performed by left-rotating the combined 32-byte quantity (**vHI:vLO**) by an amount determined by the address of the first byte of the desired data. This left-rotation is done by means of a **vperm** instruction whose control vector is generated by a Load Vector for Shift Left (**lvsl**) instruction after loading the most-significant quad word (MSQ) and least-significant quad word (LSQ) that contain the desired vector. The **lvsl** instruction uses the same address specification as the load vector indexed that loads the **vHI** component, which for big-endian ordering is the address of the desired vector.

The following instruction sequence extracts the quad word in big-endian mode:

```
lvx      vHI,rA,rB      ;# load the MSQ
lvsl     vP,rA,rB      ;# set the permute vector
addi     rB,rB,16      ;# address of LSQ
lvx      vLO,rA,rB      ;# load LSQ component
vperm    vT,vHI,vLO,vP ;# align the data
```

Note that when streaming data is used, the overhead of generating the alignment permute vector can be spread out and the latency of the loads may be covered by loop unrolling.

The process of storing a misaligned vector is essentially the reverse of that for loading; except that the code has a read-modify-write sequence. The logical algorithm is that the vector source must be right-shifted and split into two parts, each of which is merged (via a Vector Select (**vsel**) instruction) with the current contents of its MSQ and its LSQ and stored back using a store vector indexed instruction (**svx**).

The Load Vector for Shift Right (**lvsr**) instruction is used to produce the permute control vector to be used for the right-shifting. An observation is that a single register can be used for the shifted contents if a right-rotate is done. The rotate is affected by specifying the source register for both components of the Vector Permute (**vperm**); that is, a shift of a double register with the same contents in both parts results in a rotate. In addition, the same permute control vector can be used on a sequence of 1s and 0s to generate a mask for use by the **vsel** instruction to do the merging.

The complete code sequence for the store case is as follows:

```

lvx      vHI,rA,rB      ;# load current MSQ for update
lvsr     vP,rA,rB       ;# load the alignment vector
addi     rB,rB,16        ;# address of LSQ
lvx      vLO,rA,rB      ;# load the current LSQ's data
vspltib  vls,-1         ;# generate the select mask bits
vspltib  v0s,0
vperm    vMask,v0s,vls,vP ;# right rotate the select mask
vperm    vSrc,vSrc,vSrc,vP ;# right rotate the data
vsel     vLO,vSrc,vLO,vMask ;# LSQ component
vsel     vHI,vHI,vSrc,vMask ;# MSQ component
stvx     vLO,rA,rB       ;# store LSQ
addi     rB,rB,-16       ;# address of MSQ
stvx     vHI,rA,rB       ;# store MSQ

```

When fetching a linear stream of misaligned quad words using aligned quad words, the control vector need only be computed once and the time required for aligned fetches on the ends of the stream is proportioned out over that period of time. None of the data fetched internally to the stream is wasted and only gets fetched once. Thus the average time expended for a misaligned **lvx** instruction in a long sequence approaches one **lvx** and one **vperm** instruction.

### 3.1.6.2 Accessing a Misaligned Quad Word in Little-Endian Mode

The instruction sequences used to access misaligned quad-word operands in little-endian mode are similar to those used in big-endian mode. The following instruction sequence can be used to load the misaligned quad word shown in Figure 3-7 into a vector register in little-endian mode. The load alignment case is shown in Figure 3-9. The vector register **vHI** and **vLO** receive the MSQ and LSQ respectively; **vD** is the target vector register. The **lvsr** instruction uses the same address specification as the **lvx** instruction that loads **vLO**; in little-endian byte ordering this is the address of the desired misaligned quad word.

```

# Assumptions:
# Rb = / 0 and contents of Rb = 0xB
lvx      vLO,0,rB        ;# load LSQ
lvsr     vP,0,rB         ;# set permute control vector
addi     rB,rB,16        ;# address of MSQ
lvx      vHI,0,rB        ;# load MSQ
vperm    vD,vHI,vLO,vP   ;# align the data

```

Similarly, the following sequence of instructions stores the contents of register **vD** into an misaligned quad word in memory in little-endian mode.

```

# Assumptions:
# rB = / 0 and contents of rB = 0xB


lvx      vLO,0,rB        ;# load current LSQ
lvsl     vP,0,rB         ;# set permute control vector
addi     rB,rB,16        ;# address of MSQ

```

```

lvx          vHI,0,rB          # load current MSQ
vspltisb     vls,-1            # generate the select mask bits
vspltisb     v0s,0
vperm        vMask,v0s,vls,vP  # generate the select mask
vperm        vS,vS,vS,vP       # left rotate the data
vsel         vHI,vHI,vS,vMask  # insert MSQ component
vsel         vLO,vS,vLO,vMask  # insert LSQ component
stvx         vHI,0,rB          # store MSQ
addi         rB,rB,-16         # address of LSQ
stvx         vLO,0,rB          # store LSQ

```



**Figure 3-9. Little-Endian Alignment**

### 3.1.6.3 Scalar Loads and Stores

As mentioned previously, no alignment is performed for scalar load or store instructions in the AltiVec ISA. If a vector load or store address is not properly size aligned, the suitable number of least significant bits are ignored, and a size aligned transfer occurs instead. Data alignment must be performed explicitly after being brought into the registers. No assistance is provided to help align individual scalar elements that are not aligned on their natural boundary. The placement of scalar data in a vector element depends upon its address. That is, the placement of the addressed scalar is the same as if a load vector indexed instruction has been performed, except that only the addressed scalar is accessed (for cache-inhibited space); the values in the other vector elements are boundedly undefined. Also, data in the specified scalar is the same as if a store vector indexed instruction had been performed, except that only the scalar addressed is affected. No instructions are provided to assist in aligning individual scalar elements that are not aligned on their natural size boundary.

When a program knows the location of a scalar, it can perform the correct vector splats and vector permutes to move data to where it is required. For example, if a scalar is to be used as a source for a vector multiply (that is, each element multiplied by the same value), the scalar must be splatted into a vector register. Likewise, a scalar stored to an arbitrary memory location must be splatted into a vector register, and that register must be specified as the source of the store. This guarantees that the data appears in all possible positions of that scalar size for the store.

### 3.1.6.3.1 Misaligned Scalar Loads and Stores

While no direct support of misaligned scalars is provided, the load-aligning sequence for big-endian vectors described in Section 3.1.6.1, “Accessing a Misaligned Quad Word in Big-Endian Mode” can be used to position the scalar to the left vector element, which can then be used as the source for a splat. That is, the address of a scalar is also the address of the left-most element of the quad word at that address. Similarly, the read-modify-write sequences, with the mask adjusted for the scalar size, can be used to store misaligned scalars. The same is true for little-endian mode, the load-aligning sequence for little-endian vectors described Section 3.1.6.2, “Accessing a Misaligned Quad Word in Little-Endian Mode” can be used to position the scalar to the right vector element, which can then be used as the source for a splat. That is, the address of a scalar is also the address of the right-most element of the quad word at that address.

Note that while these sequences work in cache-inhibited space, the physical accesses are not guaranteed to be atomic.

### 3.1.7 Mixed-Endian Systems

In many systems, the memory model is not as simple as the examples in this chapter. In particular, big-endian systems with subordinate little-endian buses (such as PCI) comprise a mixed-endian environment.

The basic mechanism to handle this is to use the Vector Permute (**vperm**) instruction to swap bytes within data elements. The value of the permute control vector depends on the size of the elements (8, 16, 32). That is, the permute control vector performs a parallel equivalent of the PowerPC instruction Load Word Byte-Reverse Indexed (**lwbrx**) instruction, within the vector registers.

The ultimate problem is when there are misaligned, mixed-endian vectors. This can be handled by applying a vector permute of the data as required for the misaligned case, followed by the swapping vector permute on that result. Note that for streaming cases, the effect of this double permute can be accomplished by computing the swapping permute of the alignment permute vector, and then applying the resulting permute control vector to incoming data.

## 3.2 AltiVec Floating-Point —UIA

- U** There are two kinds of floating-point instructions defined for the PowerPC and AltiVec ISA—computational and noncomputational. Computational instructions consist of those operations defined by the IEEE-754 standard for 32-bit arithmetic (those that perform addition, subtraction, multiplication, and division) and the multiply-add defined by the architecture. Noncomputational floating-point instructions consist of the floating-point load and store instructions. Only the computational instructions are considered floating-point operations throughout this chapter.

The single-precision format, value representations, and computational model defined in Chapter 3, “Operand Conventions,” in *PowerPC Microprocessor Family: The Programming Environments Manual* apply to AltiVec floating-point except as follows:

- In general, no status bits are set to reflect the results of floating-point operations. The only exception is that VSCR[SAT] may be set by the Vector Convert to Fixed-Point Word instructions.
- With the exception of the two Vector Convert to Fixed-Point Word (**vctuxs**, **vctxsx**) instructions and three of the four Vector Round to Floating-Point Integer (**vrfiz**, **vrflp**, **vrflm**) instructions, all AltiVec floating-point instructions that round use the round-to-nearest rounding mode.
- Floating-point exceptions cannot cause the system error handler to be invoked.

If a function is required that is specified by the IEEE standard, is not supported by AltiVec ISA, and cannot be emulated satisfactorily using the functions that are supported by AltiVec ISA, the functions provided by the floating-point processor should be used; see Chapter 4, “Addressing Modes and Instruction Set Summary,” in *PowerPC: The Programming Environments Manual*

### 3.2.1 Floating-Point Modes

AltiVec ISA supports two floating-point modes of operation—a Java mode and a non-Java mode of operation that is useful in circumstances where real-time performance is more important than strict Java and IEEE-standard compliance.

When VSCR[NJ] is 0 (default), operations are performed in Java mode. When VSCR[NJ] is 1, operations are carried out in the non-Java mode.

#### 3.2.1.1 Java Mode

Java compliance requires compliance with only a subset of the Java/IEEE/C9X standard. The Java subset helps simplify floating-point implementations, as follows:

- Reducing the number of operations that must be supported
- Eliminating exception status flags and traps
- Producing results corresponding to all disabled exceptions thus eliminating enabling control flags
- Requiring only round-to-nearest rounding mode eliminates directed rounding modes and the associated rounding control flags.

Java compliance requires the following aspects of the IEEE standard:

- Supporting denorms as inputs and results (gradual underflow) for arithmetic operations
- Providing NaN results for invalid operations

- NaNs compare unordered with respect to everything, so that the result of any comparison of any NaN to any data type is always false

In some implementations, floating-point operations in Java mode may have somewhat longer latency on normal operands and possibly much longer latency on denormalized operands than operations in non-Java mode. This means that in Java mode overall real-time response may be somewhat worse and deadline scheduling may be subject to much larger variance than non-Java mode.

### 3.2.1.2 Non-Java Mode

In the non-Java/non-IEEE/non-C9X mode ( $VSCR[NJ] = 1$ ), gradual underflow is not performed. Instead, any instruction that would have produced a denormalized result in Java mode substitutes a correctly signed zero ( $\pm 0.0$ ) as the final result. Also, denormalized input operands are flushed to the correctly signed zero ( $\pm 0.0$ ) before being used by the instruction.

The intent of this mode is to give programmers a way to assure optimum, data-insensitive, real-time response across implementations. Another way to improve response time would be to implement denormalized operations through software emulation.

It is architecturally permitted, but strongly discouraged, for an implementation to implement only non-Java mode. In such an implementation, the  $VSCR[NJ]$  does not respond to attempts to clear it and is always read back as a 1.

No other architecturally-visible, implementation-specific deviations from this specification are permitted in either mode.

### 3.2.2 Floating-Point Infinities

Valid operations on infinities are processed according to the IEEE standard.

### 3.2.3 Floating-Point Rounding

All AltiVec floating-point arithmetic instructions use the IEEE default rounding mode, round-to-nearest. The IEEE directed rounding modes are not provided.

### 3.2.4 Floating-Point Exceptions

The following floating-point exceptions may occur during execution of AltiVec floating-point instructions.

- NaN operand exception
- Invalid operation exception
- Zero divide exception
- Log of zero exception
- Overflow exception

- Underflow exception

If an exception occurs, a result is placed into the corresponding target element as described in the following subsections. This result is the default result specified by Java, the IEEE standard, or C9X, as applicable. Recall that denormalized source values are treated as if they were zero when  $VSCR[NJ] = 1$ . The consequences regarding exceptions are as follows:

- Exceptions that can be caused by a zero source value can be caused by a denormalized source value when  $VSCR[NJ] = 1$ .
- Exceptions that can be caused by a nonzero source value cannot be caused by a denormalized source value when  $VSCR[NJ] = 1$ .

### 3.2.4.1 NaN Operand Exception

If the exponent of a floating-point number is 255 and the fraction is non-zero, then the value is a NaN. If the most significant bit of the fraction field of a NaN is zero, then the value is a signaling NaN (SNaN), otherwise it is a quiet NaN (QNaN). In all cases the sign of a NaN is irrelevant.

A NaN operand exception occurs when a source value for any of the following instructions is a NaN.

- An AltiVec instruction that would normally produce floating-point results
- Either of the two, Vector Convert to Unsigned Fixed-Point Word Saturate (**vctuxs**) or Vector Convert to Signed Fixed-Point Word Saturate (**vctxsxs**) instructions
- Any of the four vector floating-point compare instructions

The following actions are taken:

1. If the AltiVec instruction would normally produce floating-point results, the corresponding result is a source NaN selected as follows. In all cases, if the selected source NaN is an SNaN it is converted to the corresponding QNaN (by setting the high-order bit of the fraction field to 1 before being placed into the target element).
 

```

      if the element in register vA is a NaN
          then the result is that NaN
      else if the element in register vB is a NaN
          then the result is that NaN
      else if the element in register vC is a NaN
          then the result is that NaN
      
```
2. If the instruction is either of the two vector convert to fixed-point word instructions (**vctuxs**, **vctxsxs**), the corresponding result is 0x0000\_0000.  $VSCR[SAT]$  is not affected.
3. If the instruction is Vector Compare Bounds Floating-Point (**vcmpbfp[.]**), the corresponding result is 0xC000\_0000.



4. If the instruction is one of the other three vector floating-point compare instructions (**vcmpeqfp**[], **vcmpfgefp**[], **vcmpbfp**[]), the corresponding result is 0x0000\_0000.

### 3.2.4.2 Invalid Operation Exception

An invalid operation exception occurs when a source value is invalid for the specified operation. The invalid operations are as follows:

- Magnitude subtraction of infinities
- Multiplication of infinity by zero
- Vector Reciprocal Square Root Estimate Float (**vrsqrtefp**) of a negative, nonzero number or -X
- Log base 2 estimate (**vlogefp**) of a negative, nonzero number or -X

The corresponding result is the QNaN 0x7FC0\_0000. This is the single-precision format analogy of the double precision format generated QNaN described in Chapter 3, “Operand Conventions,” in *PowerPC: The Programming Environments Manual*.

### 3.2.4.3 Zero Divide Exception

A zero divide exception occurs when a Vector Reciprocal Estimate Floating-Point (**vrefp**) or Vector Reciprocal Square Root Estimate Floating-Point (**vrsqrtefp**) instruction is executed with a source value of zero.

The corresponding result is infinity, where the sign is the sign of the source value, as follows:

- $1/+0.0 \rightarrow +\infty$
- $1/-0.0 \rightarrow -\infty$
- $1/(\sqrt{+0.0}) \rightarrow +\infty$
- $1/(\sqrt{-0.0}) \rightarrow -\infty$

### 3.2.4.4 Log of Zero Exception

A log of zero exception occurs when a Vector Log Base 2 Estimate Floating-Point instruction (**vlogefp**) is executed with a source value of zero. The corresponding result is infinity. The exception cases are as follows:

- **vlogefp**  $\log_2(\pm 0.0) \rightarrow -\infty$
- **vlogefp**  $\log_2(-x) \rightarrow \text{QNaN}$ , where  $x \neq 0$

### 3.2.4.5 Overflow Exception

An overflow exception happens when either of the following conditions occur:



- For an AltiVec instruction that would normally produce floating-point results, the magnitude of what would have been the result if the exponent range were unbounded exceeds that of the largest finite single-precision number.
- For either of the two Vector Convert To Fixed-Point Word instructions (**vctuxs**, **vctxsx**), either a source value is an infinity or the product of a source value and 2 unsigned immediate value (UIMM) is a number too large to be represented in the target integer format.

The following actions are taken:

1. If the AltiVec instruction would normally produce floating-point results, the corresponding result is infinity, where the sign is the sign of the intermediate result.
2. If the instruction is Vector Convert to Unsigned Fixed-Point Word Saturate (**vctuxs**), the corresponding result is 0xFFFF\_FFFF if the source value is a positive number or +X, and is 0x0000\_0000 if the source value is a negative number or -X. VSCR[SAT] is set.
3. If the instruction is Vector Convert to Signed Fixed-Point Word Saturate (**vctfsx**), the corresponding result is 0x7FFF\_FFFF if the source value is a positive number or +X, and is 0x8000\_0000 if the source value is a negative number or -X. VSCR[SAT] is set.

### 3.2.4.6 Underflow Exception

Underflow exceptions occur only for AltiVec instructions that would normally produce floating-point results. It is detected before rounding. It occurs when a nonzero intermediate result, computed as though both the precision and the exponent range were unbounded, is less in magnitude than the smallest normalized single-precision number ( $2^{-126}$ ).

The following actions are taken:

1. If VSCR[NJ] = 0, the corresponding result is the value produced by denormalizing and rounding the intermediate result.
2. If VSCR[NJ] = 1, the corresponding result is a zero, where the sign is the sign of the intermediate result.

## 3.2.5 Floating-Point NaNs

The AltiVec floating-point data format is compliant with the Java/IEEE/C9X single-precision format. A quantity in this format can represent a signed normalized number, a signed denormalized number, a signed zero, a signed infinity, a quiet not a number (QNaN), or a signaling NaN (SNaN).

### 3.2.5.1 NaN Precedence

Whenever only one source operand of an instruction that returns a floating-point result is a NaN, then that NaN is selected as the input NaN to the instruction. When more than one source operand is a NaN, the precedence order for selecting the NaN is first from **vA** then from **vB** and then from **vC**. If the selected NaN is an SNaN, it is processed as described in

Section 3.2.5.2, “SNaN Arithmetic.” If the selected NaN is a QNaN, it is processed according to Section 3.2.5.3, “QNaN Arithmetic.”

### 3.2.5.2 SNaN Arithmetic

Whenever the input NaN to an instruction is an SNaN, a QNaN is delivered as the result, as specified by the IEEE standard when no trap occurs. The delivered QNaN is an exact copy of the original SNaN except that it is quieted; that is, the most-significant bit (msb) of the fraction is set to one (1).

### 3.2.5.3 QNaN Arithmetic

Whenever the input NaN to an instruction is a QNaN, it is propagated as the result according to the IEEE standard. All information in the QNaN is preserved through all arithmetic operations.

### 3.2.5.4 NaN Conversion to Integer

All NaNs convert to zero on conversions to integer instructions such as **vctuxs** and **vctsys**.

### 3.2.5.5 NaN Production

Whenever the result of an AltiVec operation originates a NaN (for example, an invalid operation), the NaN produced is a QNaN with the sign bit = 0, exponent field = 255, msb of the fraction field = 1, and all other bits = 0.

# Chapter 4

## Addressing Modes and Instruction Set Summary

This chapter describes instructions and addressing modes defined by the AltiVec Instruction Set Architecture (ISA), and according to the three levels of the PowerPC architecture—user instruction set architecture (UISA), virtual environment architecture (VEA), and operating environment architecture (OEA). AltiVec instructions are primarily UISA, and if otherwise they are noted in the chapter. These instructions are divided into the following categories:



- Vector integer arithmetic instructions—These include arithmetic, logical, compare, rotate and shift instructions, described in Section 4.2.1, “Vector Integer Instructions.”
- Vector floating-point arithmetic instructions—These include floating-point arithmetic instructions, as well as a discussion on floating-point modes, described in Section 4.2.2, “Vector Floating-Point Instructions.”
- Vector load and store instructions—These include load and store instructions for vector registers, described in Section 4.2.3, “Load and Store Instructions.”
- Vector permutation and formatting instructions—These include pack, unpack, merge, splat, permute, select and shift instructions, described in Section 4.2.5, “Vector Permutation and Formatting Instructions.”
- Processor control instructions—These instructions are used to read and write from the AltiVec Status and Control Register., described in Section 4.2.6, “Processor Control Instructions—UISA.”
- Memory control instructions—These instructions are used for managing of caches (user level and supervisor level), described in Section 4.3.1, “Memory Control Instructions—VEA.”

This grouping of instructions does not necessarily indicate the execution unit that processes a particular instruction or group of instructions within a processor implementation.

Integer instructions operate on byte, half-word, and word operands. Floating-point instructions operate on single-precision operands. The AltiVec ISA uses instructions that are four bytes long and word-aligned. It provides for byte, halfword, and word operand fetches and stores between memory and the vector registers (VRs).

Arithmetic and logical instructions do not read or modify memory. To use the contents of a memory location in a computation and then modify the same or another memory location, the memory contents must be loaded into a register, modified, and then written to the target location using load and store instructions.

## 4.1 Conventions

This section describes conventions used for the AltiVec instruction set. Descriptions of memory addressing, synchronization, and the AltiVec exception summary follow.

### 4.1.1 Execution Model

When used with the PowerPC instructions, AltiVec instructions can be viewed by the programmer as simply new PowerPC instructions that are freely intermixed with existing ones to provide additional features in the instruction set. PowerPC processors appear to execute instructions in program order. Some AltiVec implementations may not allow out-of-order execution and completion. Non-data dependent vector instructions may issue and execute while longer latency previously issued instructions are still in the execution stage. Register renaming is useful for AltiVec instructions to avoid stalling dispatch on false dependencies and allow maximum register name reuse in heavily unrolled loops. The execution of a sequence of instructions will not be interrupted by exceptions as the unit does not report IEEE exceptions but rather produces the default results as specified in the Java/IEEE/C9X standards. The execution of a sequence of instructions may only be interrupted by a vector load or store instruction, otherwise AltiVec instructions do not generate any exceptions.

### 4.1.2 Computation Modes

The AltiVec ISA supports the following PowerPC architecture types of implementations:

- 64-bit implementations, in that all general-purpose and floating-point registers, and some special-purpose registers (SPRs) are 64 bits long and effective addresses are 64 bits long. All 64-bit implementations have two modes of operation: the default 64-bit mode and 32-bit mode. The mode controls how an effective address is interpreted, how condition bits are set, and how the count register (CTR) is tested by branch conditional instructions.
- U** • 32-bit implementations, in that all registers except FPRs are 32 bits long and effective addresses are 32 bits long.

### 4.1.3 Classes of Instructions

AltiVec instructions follows the illegal instruction class defined by the PowerPC architecture in the section “Classes of Instructions” in Chapter 4, “Addressing Modes and Instruction Set Summary,” of the *PowerPC Microprocessor Family: The Programming*

*Environments Manual*. For AltiVec ISA, all unspecified encodings within the major opcode (04) that are not defined are illegal PowerPC instructions. The only exclusion in defining an unspecified encoding is an unused bit in an immediate field or specifier field (///).

#### 4.1.4 Memory Addressing

A program references memory using the effective (logical) address computed by the processor when it executes a load, store, or cache instruction, and when it fetches the next sequential instruction.

##### 4.1.4.1 Memory Operands

Bytes in memory are numbered consecutively starting with zero. Each number is the address of the corresponding byte.

Memory operands may be bytes, half words, words, or quad words for AltiVec instructions. The address of a memory operand is the address of its first byte (that is, of its lowest-numbered byte). Operand length is implicit for each instruction. The AltiVec ISA supports both big-endian and little-endian byte ordering. The default byte and bit ordering is big-endian; see Section 3.1.2, “AltiVec Byte Ordering,” for more information.

The natural alignment boundary of an operand of a single-register memory access instruction is equal to the operand length. In other words, the natural address of an operand is an integral multiple of the operand length. A memory operand is said to be aligned if it is aligned at its natural boundary; otherwise it is misaligned. For a detailed discussion about memory operands, see 3.1, “Data Organization in Memory and Data Transfers.”

##### 4.1.4.2 Effective Address Calculation

An effective address (EA) is the 32-bit sum computed by the processor when executing a memory access or when fetching the next sequential instruction. For a memory access instruction, if the sum of the EA and the operand length exceeds the maximum EA, the memory operand is considered to wrap around from the maximum EA through EA 0, as described in the Chapter 4, “Addressing Modes and Instruction Set Summary,” in *PowerPC Microprocessor Family: The Programming Environments Manual*.

A zero in the **rA** field indicates the absence of the corresponding address component. For the absent component, a value of zero is used for the address. This is shown in the instruction description as (**rA**|0).

In all implementations (including 32-bit mode in 64-bit implementations), the processor can modify the three low-order bits of the calculated effective address before accessing memory if the PowerPC system is operating in little-endian mode. The double words of a quad word may be swapped as well. See Section 3.1.2, “Byte Ordering,” for more information about little-endian mode.

Altivec load and store operations use register indirect with index mode and boundary align to generate effective addresses. For further details see Section 4.2.3.2.1, “Register Indirect with Index Addressing for Loads and Stores.”

## 4.2 Altivec UISA Instructions

Altivec instructions can provide additional supporting instructions to the PowerPC architecture. This section discusses the instructions defined in the Altivec user instruction set architecture (UISA).

### 4.2.1 Vector Integer Instructions

The following are categories for vector integer instructions:

- Vector integer arithmetic instructions
- Vector integer compare instructions
- Vector integer logical instructions
- Vector integer rotate and shift instructions

Integer instructions use the content of the vector registers (VRs) as source operands and place results into VRs as well. Setting the Rc bit of a vector compare instruction causes the PowerPC condition register (CR) to be updated.

The Altivec integer instructions treat source operands as signed integers unless the instruction is explicitly identified as performing an unsigned operation. For example, Vector Add Unsigned Word Modulo (**vadduwm**) and Vector Multiply Odd Unsigned Byte (**vmuloub**) instructions interpret both operands as unsigned integers.

#### 4.2.1.1 Saturation Detection

Most integer instructions have both signed and unsigned versions and many have both modulo (wrap-around) and saturating clamping modes. Saturation occurs whenever the result of a saturating instruction does not fit in the result field. Unsigned saturation clamps results to zero on underflow and to the maximum positive integer value ( $2^n-1$ , for example, 255 for byte fields) on overflow. Signed saturation clamps results to the smallest representable negative number ( $-2^{n-1}$ , for example, -128 for byte fields) on underflow, and to the largest representable positive number ( $2^{n-1}-1$ , for example, +127 for byte fields) on overflow. When a modulo instruction is used, the resultant number truncates overflow or underflow for the length (byte, half word, word, quad word) and type of operand (unsigned, signed). The Altivec ISA provides a way to detect saturation and sets the SAT bit in the Vector Status and Control Register (VSCR[SAT]) in a saturating instruction.

Borderline cases that generate results equal to saturation values, for example unsigned  $0+0 \rightarrow 0$  and unsigned byte  $1+254 \rightarrow 255$ , are not considered saturation conditions and do not cause VSCR[SAT] to be set.

The VSCR[SAT] can be set by the following types of integer, floating-point, and formatting instructions:

- Move to VSCR (**mtvscr**)
- Vector add integer with saturation (**vaddubs**, **vadduhs**, **vadduws**, **vaddsbs**, **vaddshs**, **vaddsws**)
- Vector subtract integer with saturation (**vsububs**, **vsubuhs**, **vsubuws**, **vsubsbbs**, **vsubshs**, **vsubsws**)
- Vector multiply-add integer with saturation (**vmhaddshs**, **vmhraddshs**)
- Vector multiply-sum with saturation (**vmsumuhs**, **vmsumshs**, **vsumsws**)
- Vector sum-across with saturation (**vsumsws**, **vsum2sws**, **vsum4sbs**, **vsum4shs**, **vsum4ubs**)
- Vector pack with saturation (**vpkuhus**, **vpkuwus**, **vpkshus**, **vpkswus**, **vpkshss**, **vpkswss**)
- Vector convert to fixed-point with saturation (**vctuhs**, **vctshs**)

Note that only instructions that explicitly call for saturation can set VSCR[SAT]. Modulo integer instructions and floating-point arithmetic instructions never set VSCR[SAT]. For further details see Section 2.1.1, “The Vector Status and Control Register (VSCR).”

#### 4.2.1.2 Vector Integer Arithmetic Instructions

Table 4-1 lists the integer arithmetic instructions for the PowerPC processors.

**Table 4-1. Vector Integer Arithmetic Instructions**

Name	Mnemonic	Syntax	Operation
Vector Add Unsigned Integer [b,h,w] Modulo	<b>vaddubm</b> <b>vadduhm</b> <b>vadduwm</b>	<b>vD,vA,vB</b>	Place the sum (vA[unsigned integer elements]) + (vB[unsigned integer elements]) into vD[unsigned integer elements] using modulo arithmetic.  For <b>b</b> , byte, integer length = 8 bits = 1 byte, add 16 unsigned integers from vA to the corresponding 16 unsigned integers from vB  For <b>h</b> , half word, integer length = 16 bits = 2 bytes, add 8 unsigned integers from vA to the corresponding 8 unsigned integers from vB  For <b>w</b> , word, integer length = 32 bits = 4 bytes, add 4 unsigned integers from vA to the corresponding 4 unsigned integers from vB  Note: unsigned or signed integers can be used with these instructions



**Table 4-1. Vector Integer Arithmetic Instructions (Continued)**

Name	Mnemonic	Syntax	Operation
Vector Add Unsigned Integer [b,h,w] Saturate	<b>vaddubs</b> <b>vadduhs</b> <b>vadduws</b>	<b>vD,vA,vB</b>	<p>Place the sum (<b>vA</b>[unsigned integer elements]) + (<b>vB</b>[unsigned integer elements]) into <b>vD</b>[unsigned integer elements] using saturate clamping mode. Saturate clamping mode means if the resulting sum is <math>&gt;(2^n-1)</math> saturate to <math>(2^n-1)</math>, where <math>n = \mathbf{b,h,w}</math>.</p> <p>For <b>b</b>, byte, integer length = 8 bits = 1 byte, add 16 unsigned integers from <b>vA</b> to the corresponding 16 unsigned integers from <b>vB</b></p> <p>For <b>h</b>, half word, integer length = 16 bits = 2 bytes, add 8 unsigned integers from <b>vA</b> to the corresponding 8 unsigned integers from <b>vB</b></p> <p>For <b>w</b>, word, integer length = 32 bits = 4 bytes, add 4 unsigned integers from <b>vA</b> to the corresponding 4 unsigned integers from <b>vB</b></p> <p>If the result saturates, VSCR[SAT] is set.</p>
Vector Add Signed Integer[b,h,w] Saturate	<b>vaddsbs</b> <b>vaddshs</b> <b>vddsws</b>	<b>vD,vA,vB</b>	<p>Place the sum (<b>vA</b>[signed integer elements]) + (<b>vB</b>[signed integer elements]) into <b>vD</b>[signed integer elements] using saturate clamping mode. Saturate clamping mode means:</p> <p>if the sum is <math>&gt;(2^{n-1}-1)</math> saturate to <math>(2^{n-1}-1)</math> and</p> <p>if <math>&lt; (-2^{n-1}-1)</math> saturate to <math>(-2^{n-1}-1)</math>, where <math>n = \mathbf{b,h,w}</math>.</p> <p>For <b>b</b>, byte, integer length = 8 bits = byte, add 16 signed integers from <b>vA</b> to the corresponding 16 signed integers from <b>vB</b></p> <p>For <b>h</b>, half word, integer length = 16 bits = 2 bytes, add 8 signed integers from <b>vA</b> to the corresponding 8 signed integers from <b>vB</b></p> <p>For <b>w</b>, word, integer length = 32 bits = 4 bytes, add 4 signed integers from <b>vA</b> to the corresponding 4 signed integers from <b>vB</b></p> <p>If the result saturates, VSCR[SAT] is set.</p>
Vector Add and Write Carry-out Unsigned Word	<b>vaddcuw</b>	<b>vD,vA,vB</b>	<p>Take the carry out of summing (<b>vA</b>) + (<b>vB</b>) and place it into <b>vD</b>.</p> <p>For <b>w</b>, word, integer length = 32 bits = 2 bytes, add 4 unsigned integers from <b>vA</b> to the corresponding 4 unsigned integers from <b>vB</b> and the resulting carry outs are correspondingly placed in <b>vD</b>.</p>
Vector Subtract Unsigned Integer Modulo	<b>vsububm</b> <b>vsubuhm</b> <b>vsubuwm</b>	<b>vD,vA,vB</b>	<p>Place the unsigned integer sum (<b>vA</b>) - (<b>vB</b>) into <b>vD</b> using modulo arithmetic.</p> <p>For <b>b</b>, byte, integer length = 8 bits = 1 byte, subtract 16 unsigned integers in <b>vB</b> from the corresponding 16 unsigned integers in <b>vA</b></p> <p>For <b>h</b>, half word, integer length = 16 bits = 2 bytes, subtract 8 unsigned integers in <b>vB</b> from the corresponding 8 unsigned integers in <b>vA</b></p> <p>For <b>w</b>, word, integer length = 32 bits = 4 bytes, subtract 4 unsigned integers in <b>vB</b> from the corresponding 4 unsigned integers in <b>vA</b></p> <p>Note that unsigned or signed integers can be used with these instructions</p>



**Table 4-1. Vector Integer Arithmetic Instructions (Continued)**

Name	Mnemonic	Syntax	Operation
Vector Subtract Unsigned Integer Saturate	<b>vsububs</b> <b>vsubuhs</b> <b>vsubuws</b>	<b>vD,vA,vB</b>	<p>Place the unsigned integer sum <b>vA</b> - <b>vB</b> into <b>vD</b> using saturate clamping mode, that is, if the sum &lt; 0, it saturates to 0 corresponding to <b>b,h,w</b>.</p> <p>For <b>b</b>, byte, integer length = 8 bits = 1 byte, subtract 16 unsigned integers in <b>vB</b> from the corresponding 16 unsigned integers in <b>vA</b></p> <p>For <b>h</b>, half word, integer length = 16 bits = 2 bytes, subtract 8 unsigned integers in <b>vB</b> from the corresponding 8 unsigned integers in <b>vA</b></p> <p>For <b>w</b>, word, integer length = 32 bits = 4 bytes, subtract 4 unsigned integers in <b>vB</b> from the corresponding 4 unsigned integers in <b>vA</b></p> <p>If the result saturates, VSCR[SAT] is set.</p>
Vector Subtract Signed Integer Saturate	<b>vsububs</b> <b>vsubuhs</b> <b>vsubuws</b>	<b>vD,vA,vB</b>	<p>Place the signed integer sum (<b>vA</b>) - (<b>vB</b>) into <b>vD</b> using saturate clamping mode. Saturate clamping mode means:</p> <p>if the sum is &gt; (<math>2^{n-1}-1</math>) saturate to (<math>2^{n-1}-1</math>) and</p> <p>if &lt; (<math>-2^{n-1}-1</math>) saturate to (<math>-2^{n-1}-1</math>), where <math>n = \mathbf{b,h,w}</math>.</p> <p>For <b>b</b>, byte, integer length = 8 bits = 1 byte, subtract 16 signed integers in <b>vB</b> from the corresponding 16 signed integers in <b>vA</b></p> <p>For <b>h</b>, half word, integer length = 16 bits = 2 bytes, subtract 8 signed integers in <b>vB</b> from the corresponding 8 signed integers in <b>vA</b></p> <p>For <b>w</b>, word, integer length = 32 bits = 4 bytes, subtract 4 signed integers in <b>vB</b> from the corresponding 4 signed integers in <b>vA</b></p>
Vector Subtract and Write Carry-out Unsigned Word	<b>vasubcuw</b>	<b>vD,vA,vB</b>	<p>Take the carry out of the sum (<b>vA</b>) - (<b>vB</b>) and place it into <b>vD</b>.</p> <p>For <b>w</b>, word, integer length = 32 bits = 2 bytes, subtract 4 unsigned integers in <b>vB</b> from the corresponding 4 unsigned integers in <b>vA</b> and place the resulting carry outs into <b>vD</b>.</p>
Vector Multiply Odd Unsigned Integer [b,h] Modulo	<b>vmuloub</b> <b>vmulouh</b>	<b>vD,vA,vB</b>	<p>Place the unsigned integer products of (<b>vA</b>) * (<b>vB</b>) into <b>vD</b> using modulo arithmetic mode.</p> <p>For <b>b</b>, byte, integer length = 8 bits = 1 byte, multiply 8 odd-numbered unsigned integer byte elements from <b>vA</b> to the corresponding 8 odd-numbered unsigned integer byte elements from <b>vB</b> resulting in 8 unsigned integer half-word products in <b>vD</b>.</p> <p>For <b>h</b>, half word, integer length = 16 bits = 2 bytes, multiply 4 odd-numbered unsigned integer half word elements from <b>vA</b> to the corresponding 4 odd numbered unsigned integer half-word elements from <b>vB</b> resulting in 4 unsigned integer word products in <b>vD</b>.</p>
Vector Multiply Odd Signed Integer [b,h] Modulo	<b>vmulosb</b> <b>vmulosh</b>	<b>vD,vA,vB</b>	<p>Place the signed integer product of (<b>vA</b>) * (<b>vB</b>) into <b>vD</b> using modulo arithmetic mode.</p> <p>For <b>b</b>, byte, integer length = 8 bits = 1 byte, multiply 8 odd-numbered signed integer byte elements from <b>vA</b> to 8 odd-numbered signed integer byte elements from <b>vB</b> resulting in 8 signed integer half-word products in <b>vD</b>.</p> <p>For <b>h</b>, half word, integer length = 16 bits = 2 bytes, multiply 4 odd-numbered signed integer half word elements from <b>vA</b> to 4 odd-numbered signed integer half word elements from <b>vB</b> resulting in 4 signed integer word products in <b>vD</b>.</p>

**Table 4-1. Vector Integer Arithmetic Instructions (Continued)**

Name	Mnemonic	Syntax	Operation
Vector Multiply Even Unsigned Integer [b,h] Modulo	<b>vmuleub</b> <b>vmuleuh</b>	<b>vD,vA,vB</b>	<p>Place the unsigned integer products of (<b>vA</b>) * (<b>vB</b>) into <b>vD</b> using modulo arithmetic mode.</p> <p>For <b>b</b>, byte, integer length = 8 bits = 1 byte, multiply 8 even-numbered unsigned integer byte elements from <b>vA</b> to 8 even-numbered unsigned integer byte elements from <b>vB</b> resulting in 8 unsigned integer half-word products in <b>vD</b>.</p> <p>For <b>h</b>, half word, integer length = 16 bits = 2 bytes, multiply 4 even-numbered unsigned integer half-word elements from <b>vA</b> to 4 even-numbered unsigned integer half-word elements from <b>vB</b> resulting in 4 unsigned integer word products in <b>vD</b>.</p>
Vector Multiply Even Signed Integer [b,h] Modulo	<b>vmulesb</b> <b>vmulesh</b>	<b>vD,vA,vB</b>	<p>Place the signed integer product of (<b>vA</b>) * (<b>vB</b>) into <b>vD</b> using modulo arithmetic mode.</p> <p>For <b>b</b>, byte, integer length = 8 bits = 1 byte, multiply 8 even-numbered signed integer byte elements from <b>vA</b> to 8 even-numbered signed integer byte elements from <b>vB</b> resulting in 8 signed integer half-word products in <b>vD</b>.</p> <p>For <b>h</b>, half word, integer length = 16 bits = 2 bytes, multiply 4 even-numbered signed integer half-word elements from <b>vA</b> to 4 even-numbered signed integer half-word elements from <b>vB</b> resulting in 4 signed integer word products in <b>vD</b>.</p>
Vector Multiply-High and Add Signed Half-Word Saturate	<b>vmhaddshs</b>	<b>vD,vA,vB,vC</b>	<p>The 17 most significant bits (msb's) of the product of (<b>vA</b>) * (<b>vB</b>) adds to sign-extended <b>vC</b> and places the result into <b>vD</b>.</p> <p>For <b>h</b>, half word, integer length = 16 bits = 2 bytes, multiply the 8 signed half words from <b>vA</b> with the corresponding 8 signed half words from <b>vB</b> to produce a 32-bit intermediate product and then take the 17 msb's (bits 0–16) of the 8 intermediate products and add them to the 8 sign-extended half words in <b>vC</b>, place the 8 half-word saturated results in <b>vD</b>. If the intermediate product is as follows:</p> <p>&gt; <math>(2^{15}-1)</math> saturate to <math>(2^{15}-1)</math> and if</p> <p>&lt; <math>-2^{15}</math> saturate to <math>-2^{15}</math>.</p> <p>If the results saturates, VSCR[SAT] is set.</p>
Vector Multiply-High Round and Add Signed Half-Word Saturate	<b>vmhraddshs</b>	<b>vD,vA,vB,vC</b>	<p>Add the rounded product of (<b>vA</b>) * (<b>vB</b>) to sign-extended <b>vC</b> and place the result into <b>vD</b>.</p> <p>For <b>h</b>, half word, integer length = 16 bits = 2 bytes, multiply the 8 signed integers from <b>vA</b> to the corresponding 8 signed integers from <b>vB</b> and then round the 8 immediate products by adding the value 0x0000_4000 to it. Then add the most significant bits (msb's), bits 0–16, of the 8 rounded immediate products to the 8 sign-extended values in <b>vC</b> and place the 8 signed half-word saturated results into <b>vD</b>. If the intermediate product is:</p> <p>&gt; <math>(2^{15}-1)</math> saturate to <math>(2^{15}-1)</math> and if</p> <p>&lt; <math>-2^{15}</math> saturate to <math>-2^{15}</math>.</p> <p>If the result saturates, VSCR[SAT] is set.</p>

**Table 4-1. Vector Integer Arithmetic Instructions (Continued)**

Name	Mnemonic	Syntax	Operation
Vector Multiply-Low and Add Unsigned Half-Word Modulo	<b>vmladduhm</b>	<b>vD,vA,vB,vC</b>	<p>Add the product of (vA) * (vB) to zero-extended vC and place into vD.</p> <p>For <b>h</b>, half word, integer length = 16 bits = 2 bytes, multiply the 8 signed integers from vA to the corresponding 8 signed integers from vB to produce a 32-bit intermediate product. The 16-bit value in vC is zero-extended to 32 bits and added to the intermediate product and the lower 16 bits of the sum (bit 16–31) is placed in vD.</p> <p>Note that unsigned or signed integers can be used with these instructions</p>
Vector Multiply-Sum Unsigned Integer [b,h] Modulo	<b>vmsumubm</b> <b>vmsumuhm</b>	<b>vD,vA,vB,vC</b>	<p>The product of (vA) * (vB) is added to zero-extended vC and placed into vD using modulo arithmetic.</p> <p>For <b>b</b>, byte, integer length = 8 bits = 1 byte, multiply 4 unsigned integer bytes from a word element in vA by the corresponding 4 unsigned integer bytes in a word element in vB and the sum of these products are added to the zero-extended unsigned integer word element in vC and then placed the unsigned integer word result into vD, following this process for each 4-word element in vA and vB.</p> <p>For <b>h</b>, half word, integer length = 16 bits = 2 bytes, multiply 2 unsigned integer half words from a word element in vA by the corresponding 2 unsigned integer half words in a word element in vB and the sum of these products are added to zero-extended unsigned integer word element in vC and then place the unsigned integer word result into vD, following this process for each 4 word element in vA and vB.</p>
Vector Multiply-Sum Signed Half-Word Saturate	<b>vmsumshs</b>	<b>vD,vA,vB,vC</b>	<p>Add the product of (vA) * (vB) to vC and place the result into vD using saturate clamping mode.</p> <p>For <b>h</b>, half word, integer length = 16 bits = 2 bytes, multiply 2 signed integer half words from a word element in vA by the corresponding 2 signed integer half words in a word element in vB. Add the sum of these products to the signed integer word element in vC and then place the signed integer word result into vD, (following this process for each 4-word element in vA and vB). If the intermediate result is <math>&gt; (2^{31}-1)</math>, saturate to <math>(2^{31}-1)</math> and if the result is <math>&lt; -2^{31}</math>, saturate to <math>-2^{31}</math>.</p> <p>If the result saturates, VSCR[SAT] is set.</p>
Vector Multiply-Sum Unsigned Half-Word Saturate	<b>vmsumuhs</b>	<b>vD,vA,vB,vC</b>	<p>Add the product of (vA) * (vB) to zero-extended vC and place the result into vD using saturate clamping mode.</p> <p>For <b>h</b>, half word, integer length = 16 bits = 2 bytes, multiply 2 unsigned integer half words from a word element in vA by the corresponding 2 unsigned integer half words in a word element in vB. Add the sum of these products to the zero-extended unsigned integer word element in vC and then place the unsigned integer word result into vD, (following this process for each 4-word element in vA and vB). If the intermediate result is <math>&gt; (2^{32}-1)</math> saturate to <math>(2^{32}-1)</math>.</p> <p>If the result saturates, VSCR[SAT] is set.</p>

**Table 4-1. Vector Integer Arithmetic Instructions (Continued)**

Name	Mnemonic	Syntax	Operation
Vector Multiply-Sum Mixed Byte Modulo	<b>vmsumshm</b>	<b>vD,vA,vB,vC</b>	<p>Add the product of (<b>vA</b>) * (<b>vB</b>) to <b>vC</b> and place into <b>vD</b> using modulo arithmetic.</p> <p>For <b>b</b>, byte, integer length = 8 bits = 1 byte, multiply 4 signed integer bytes from a word element in <b>vA</b> by the corresponding 4 unsigned integer bytes from a word element in <b>vB</b>. Add the sum of these 4 signed products to the signed integer word element in <b>vC</b> and then place the signed integer word result into <b>vD</b>, following this process for each 4-word element in <b>vA</b> and <b>vB</b>.</p>
Vector Multiply-Sum Signed Half-Word Modulo	<b>vmsumshw</b>	<b>vD,vA,vB,vC</b>	<p>Add the product of (<b>vA</b>) * (<b>vB</b>) to <b>vC</b> and place into <b>vD</b> using modulo arithmetic.</p> <p>For <b>h</b>, half word, integer length = 16 bits = 2 bytes, multiply 2 signed integer half words from a word element in <b>vA</b> by the corresponding 2 signed integer half words in a word element in <b>vB</b>. Add the sum of these 2 products to the signed integer word element in <b>vC</b> and then place the signed integer word result into <b>vD</b>, following this process for each 4-word element in <b>vA</b> and <b>vB</b>.</p>
Vector Sum Across Signed Word Saturate	<b>vmsumsws</b>	<b>vD,vA,vB</b>	<p>Place the sum of signed word elements in <b>vA</b> and the word in <b>vB</b>[96–127] into <b>vD</b>.</p> <p>For <b>w</b>, word, integer length = 32 bits = 4 bytes, add the sum of the 4 signed integer word elements in <b>vA</b> to the word element in <b>vB</b>[96–127]. If the intermediate product is <math>&gt; (2^{31}-1)</math> saturate to <math>(2^{31}-1)</math> and if <math>&lt; -2^{31}</math> saturate to <math>-2^{31}</math>. Place the signed integer result in <b>vD</b>[96–127], <b>vD</b>[0–95] are cleared.</p>
Vector Sum Across Partial (1/2) Signed Word Saturate	<b>vmsum2sws</b>	<b>vD,vA,vB</b>	<p>Add <b>vA</b>[word 0 + word 1] + <b>vB</b>[word 1] and place in <b>vD</b>[word 1]. Repeat only add <b>vA</b>[word 2 + word 3] + <b>vB</b>[word 3] and place in <b>vD</b>[word 3].</p> <p>word 0 = Bits 0-31 word 1 = Bits 32-63 word 2 = Bits 64-95 word 3 = Bits 96-127,</p> <p>See Figure1-2, “Big-Endian Byte Ordering for a Vector Register” for a picture of what the word elements would look like in a vector register.</p> <p>Add the sum of word 0 and word 1 of <b>vA</b> to word 1 of <b>vB</b> using saturate clamping mode and place the result is into word 1of <b>vD</b>. Then add the sum of word 2 and word 3 of (<b>vA</b>) to word 3 of <b>vB</b> using saturate clamping mode and place those results into word 3 in <b>vD</b>. If the intermediate result for either calculation is <math>&gt; (2^{31}-1)</math> then saturate to <math>(2^{31}-1)</math> and if <math>&lt; -2^{31}</math> then saturate to <math>-2^{31}</math>.</p> <p>If the result saturates, VSCR[SAT] is set.</p>
Vector Sum Across Partial (1/4) Unsigned Byte Saturate	<b>vmsum4ubs</b>	<b>vD,vA,vB</b>	<p>Add <b>vA</b>[sum of 4 byte elements in word] and <b>vB</b>[word element] then place in <b>vD</b>[word element] using saturate clamping mode.</p> <p>For <b>b</b>, byte, integer length = 8 bits = 1 byte, for each word element in <b>vB</b>, add the sum of 4 unsigned bytes in the word in <b>vA</b> to the unsigned word element in <b>vB</b> and then place the results into the corresponding unsigned word element in <b>vD</b>. If the intermediate result for is <math>&gt; (2^{32}-1)</math> it saturates to <math>(2^{32}-1)</math>.</p> <p>If the result saturates, VSCR[SAT] is set.</p>

**Table 4-1. Vector Integer Arithmetic Instructions (Continued)**

Name	Mnemonic	Syntax	Operation
Vector Sum Across Partial (1/4) Signed Integer Saturate	<b>vmsum4sbs</b> <b>vmsum4shs</b>	<b>vD,vA,vB</b>	<p>Add <b>vA</b>[sum of signed integer elements in word] and <b>vB</b>[word element] then place in <b>vD</b>[word element] using saturate clamping mode.</p> <p>For <b>b</b>, byte, integer length = 8 bits = 1 byte, for each word element in <b>vB</b>, add the sum of 4 signed bytes in the word in <b>vA</b> to the signed word element in <b>vB</b> and then place the results into the corresponding signed word element in <b>vD</b>. If the intermediate result is <math>&gt; (2^{31}-1)</math> then saturate to <math>(2^{31}-1)</math> and if <math>&lt; -2^{31}</math> then saturate to <math>-2^{31}</math>.</p> <p>For <b>h</b>, half word, integer length = 16 bits = 2 bytes, for each word element in <b>vB</b>, add the sum of 2 signed half words in the word in <b>vA</b> to the signed word element in <b>vB</b> and then place the results into the corresponding signed word element in <b>vD</b>. If the intermediate result is <math>&gt; (2^{31}-1)</math> then saturate to <math>(2^{31}-1)</math> and if <math>&lt; -2^{31}</math> then saturate to <math>-2^{31}</math>.</p> <p>If the result saturates, VSCR[SAT] is set.</p>
Vector Average Unsigned Integer	<b>vavgub</b> <b>vavguh</b> <b>vavguw</b>	<b>vD,vA,vB</b>	<p>Add the sum of (<b>vA</b>[unsigned integer elements]+ <b>vB</b>[unsigned integer elements]) +1 and place into <b>vD</b> using modulo arithmetic.</p> <p>For <b>b</b>, byte, integer length = 8 bits = 1 byte, add 16 unsigned integers from <b>vA</b> to 16 unsigned integers from <b>vB</b> and then add 1 to the sums and place the high order result in <b>vD</b>.</p> <p>For <b>h</b>, half word, integer length = 16 bits = 2 bytes, add 8 unsigned integers from <b>vA</b> to 8 unsigned integers from <b>vB</b> and then add 1 to the sums and place the high order result in <b>vD</b>.</p> <p>For <b>w</b>, word, integer length = 32 bits = 4 bytes, add 4 unsigned integers from <b>vA</b> to 4 unsigned integers from <b>vB</b> and then add 1 to the sums and place the high order result in <b>vD</b>.</p> <p>If the result saturates, VSCR[SAT] is set.</p>
Vector Average Signed Integer	<b>vavgsb</b> <b>vavgsh</b> <b>vavgsw</b>	<b>vD,vA,vB</b>	<p>Add the sum of (<b>vA</b>[signed integer elements]+ <b>vB</b>[signed integer elements]) +1 and place into <b>vD</b> using modulo arithmetic.</p> <p>For <b>b</b>, byte, integer length = 8 bits = 1 byte, add 16 signed integers from <b>vA</b> to 16 signed integers from <b>vB</b> and then add 1 to the sums and place the high order result in <b>vD</b>.</p> <p>For <b>h</b>, half word, integer length = 16 bits = 2 bytes, add 8 signed integers from <b>vA</b> to 8 signed integers from <b>vB</b> and then add 1 to the sums and place the high order result in <b>vD</b>.</p> <p>For <b>w</b>, word, integer length = 32 bits = 4 bytes, add 4 signed integers from <b>vA</b> to 4 signed integers from <b>vB</b> and then add 1 to the sums and place the high order result in <b>vD</b>.</p>
Vector Maximum Unsigned Integer	<b>vmaxub</b> <b>vmaxuh</b> <b>vmaxw</b>	<b>vD,vA,vB</b>	<p>Compare the maximum of <b>vA</b> and <b>vB</b> unsigned integers for each integer value and which ever value is larger, place that unsigned integer value into <b>vD</b></p> <p>For <b>b</b>, byte, integer length = 8 bits = 1 byte, compare 16 unsigned integers from <b>vA</b> with 16 unsigned integers from <b>vB</b>.</p> <p>For <b>h</b>, half word, integer length = 16 bits = 2 bytes, compare 8 unsigned integers from <b>vA</b> with 8 unsigned integers from <b>vB</b>.</p> <p>For <b>w</b>, word, integer length = 32 bits = 4 bytes, compare 4 unsigned integers from <b>vA</b> with 4 unsigned integers from <b>vB</b>.</p>

**Table 4-1. Vector Integer Arithmetic Instructions (Continued)**

Name	Mnemonic	Syntax	Operation
Vector Maximum Signed Integer	<b>vmaxsb</b> <b>vmaxsh</b> <b>vmaxsw</b>	<b>vD,vA,vB</b>	<p>Compare the maximum of <b>vA</b> and <b>vB</b> signed integers for each integer value and which ever value is larger, place that signed integer value into <b>vD</b></p> <p>For <b>b</b>, byte, integer length = 8 bits =1 byte, compare 16 signed integers from <b>vA</b> with 16 signed integers from <b>vB</b></p> <p>For <b>h</b>, half word, integer length =16 bits = 2 bytes, compare 8 signed integers from <b>vA</b> with 8 signed integers from <b>vB</b></p> <p>For <b>w</b>, word, integer length = 32 bits = 4 bytes, compare 4 signed integers from <b>vA</b> with 4 signed integers from <b>vB</b></p>
Vector Minimum Unsigned Integer	<b>vminub</b> <b>vminuh</b> <b>vminw</b>	<b>vD,vA,vB</b>	<p>Compare the minimum of <b>vA</b> and <b>vB</b> unsigned integers for each integer value and which ever value is smaller, place that unsigned integer value into <b>vD</b></p> <p>For <b>b</b>, byte, integer length = 8 bits = 1 byte, compare 16 unsigned integers from <b>vA</b> with 16 unsigned integers from <b>vB</b></p> <p>For <b>h</b>, half word, integer length = 16 bits = 2 bytes, compare 8 unsigned integers from <b>vA</b> with 8 unsigned integers from <b>vB</b></p> <p>For <b>w</b>, word, integer length = 32 bits = 4 bytes, compare 4 unsigned integers from <b>vA</b> with 4 unsigned integers from <b>vB</b></p>
Vector Minimum Signed Integer	<b>vminsb</b> <b>vminsh</b> <b>vminsw</b>	<b>vD,vA,vB</b>	<p>Compare the minimum of <b>vA</b> and <b>vB</b> signed integers for each integer value and which ever value is smaller, place that signed integer value into <b>vD</b>.</p> <p>For <b>b</b>, byte, integer length = 8 bits = 1 byte, compare 16 signed integers from <b>vA</b> with 16 signed integers from <b>vB</b></p> <p>For <b>h</b>, half word, integer length = 16 bits = 2 bytes, compare 8 signed integers from <b>vA</b> with 8 signed integers from <b>vB</b></p> <p>For <b>w</b>, word, integer length = 32 bits = 4 bytes, compare 4 signed integers from <b>vA</b> with 4 signed integers from <b>vB</b></p>

### 4.2.1.3 Vector Integer Compare Instructions

The vector integer compare instructions algebraically or logically compare the contents of the elements in vector register **vA** with the contents of the elements in **vB**. Each compare result vector is comprised of TRUE (0xFF, 0xFFFF, 0xFFFFFFFF) or FALSE (0x00, 0x0000, 0x00000000) elements of the size specified by the compare source operand element (byte, half word, or word). The result vector can be directed to any vector register and can be manipulated with any of the instructions as normal data, for example, combining condition results.

Vector compares provide equal-to and greater-than predicates. Others are synthesized from these by logically combining and/or inverting result vectors.

If the record bit (Rc) is set in the integer compare instructions (shown in Table 4-3) it can optionally set the CR6 field of the PowerPC condition register. If Rc = 1 in the vector

integer compare instruction, then CR6 is set to reflect the result of the comparison, as follows in Table 4-2.

**Table 4-2. CR6 Field Bit Settings for Vector Integer Compare Instructions**

CR Bit	CR6 Bit	Vector Compare
24	0	1 Relation is true for all element pairs (that is, vD is set to all ones)
25	1	0
26	2	1 Relation is false for all element pairs (that is, register vD is cleared)
27	3	0

Table 4-3 summarizes the vector integer compare instructions.,

**Table 4-3. Vector Integer Compare Instructions**

Name	Mnemonic	Syntax	Operation
Vector Compare Greater than Unsigned Integer	<b>vcmpgtub[.]</b> <b>vcmpgtuh[.]</b> <b>vcmpgtuw[.]</b>	<b>CR06,vD,vA,vB</b>	<p>Compare the value in <b>vA</b> with the value in <b>vB</b>, treating the operands as unsigned integers. Place the result of the comparison into the <b>vD</b> field specified by operand <b>vD</b>.</p> <p>if <b>vA &gt; vB</b> then <b>vD = 1's</b>; otherwise <b>vD = 0's</b></p> <p>If the record bit (<b>Rc</b>) is set in the vector compare instruction then</p> <p><b>vD == 1's</b>, (all elements true) then <b>CR6[0]</b> is set</p> <p><b>vD == 0's</b>, (all elements false) then <b>CR6[2]</b> is set</p> <p>For <b>b</b>, byte, integer length = 8 bits = 1 byte, compare 16 unsigned integers from <b>vA</b> to 16 unsigned integers from <b>vB</b> and place the results in the corresponding 16 elements in <b>vD</b></p> <p>For <b>h</b>, half word, integer length = 16 bits = 2 bytes, compare 8 unsigned integers from <b>vA</b> to 8 unsigned integers from <b>vB</b> and place the results in the corresponding 8 elements in <b>vD</b></p> <p>For <b>w</b>, word, integer length = 32 bits = 4 bytes, compare 4 unsigned integers from <b>vA</b> to 4 unsigned integers from <b>vB</b> and place the results in the corresponding 4 elements in <b>vD</b>.</p>



**Table 4-3. Vector Integer Compare Instructions (Continued)**

Name	Mnemonic	Syntax	Operation
Vector Compare Greater Than Signed Integer	<b>vcmpgtsb</b> [.] <b>vcmpgtsh</b> [.] <b>vcmpgtsw</b> [.]	<b>CR06,vD,vA,vB</b>	<p>Compare the value in <b>vA</b> with the value in <b>vB</b>, treating the operands as signed integers. Place the result of the comparison into the <b>vD</b> field specified by operand <b>vD</b></p> <p>if <b>vA</b> &gt; <b>vB</b> then <b>vD</b> = 1's; otherwise <b>vD</b> = 0's</p> <p>If the record bit (Rc) is set in the vector compare instruction then <b>vD</b> == 1's, (all elements true) then CR6[0] is set <b>vD</b> == 0's, (all elements false) then CR6[2] is set</p> <p>For <b>b</b>, byte, integer length = 8 bits = 1 byte, compare 16 signed integers from <b>vA</b> to 16 signed integers from <b>vB</b> and place the results in the 16 corresponding elements in <b>vD</b></p> <p>For <b>h</b>, half word, integer length = 16 bits = 2 bytes, compare 8 signed integers from <b>vA</b> to 8 signed integers from <b>vB</b> and place the results in the 8 corresponding elements in <b>vD</b></p> <p>For <b>w</b>, word, integer length = 32 bits = 4 bytes, compare 4 signed integers from <b>vA</b> to 4 signed integers from <b>vB</b> and place the results in the 4 corresponding elements in <b>vD</b></p>
Vector Compare Equal To Unsigned Integer	<b>vcmpequb</b> [.] <b>vcmpequh</b> [.] <b>vcmpequw</b> [.]	<b>vD,vA,vB</b>	<p>Compare the value in <b>vA</b> with the value in <b>vB</b>, treating the operands as unsigned integers. Place the result of the comparison into the <b>vD</b> field specified by operand <b>vD</b>.</p> <p>if <b>vA</b> = <b>vB</b> then <b>vD</b> = 1's; otherwise <b>vD</b> = 0's</p> <p>If the record bit (Rc) is set in the vector compare instruction then <b>vD</b> == 1's, (all elements true) then CR6[0] is set <b>vD</b> == 0's, (all elements false) then CR6[2] is set</p> <p>For <b>b</b>, byte, integer length = 8 bits = 1 byte, compare 16 unsigned integers from <b>vA</b> to 16 unsigned integers from <b>vB</b> and place the results in the corresponding 16 elements in <b>vD</b></p> <p>For <b>h</b>, half word, integer length = 16 bits = 2 bytes, compare 8 unsigned integers from <b>vA</b> to 8 unsigned integers from <b>vB</b> and place the results in the corresponding 8 elements in <b>vD</b></p> <p>For <b>w</b>, word, integer length = 32 bits = 4 bytes, compare 4 unsigned integers from <b>vA</b> to 4 unsigned integers from <b>vB</b> and place the results in the corresponding 4 elements in <b>vD</b>.</p> <p>Note: <b>vcmpequb</b>[], <b>vcmpequh</b>[], and <b>vcmpequw</b>[] can use both unsigned and signed integers</p>

#### 4.2.1.4 Vector Integer Logical Instructions

The vector integer logical instructions shown in . Vector Integer Logical Instructions on page 14 perform bit-parallel operations on the operands.

**Table 4-4. Vector Integer Logical Instructions**

Name	Mnemonic	Syntax	Operation
Vector Logical AND	<b>vand</b>	<b>vD,vA,vB</b>	AND the contents of <b>vA</b> with <b>vB</b> and place the result into <b>vD</b> .
Vector Logical OR	<b>vor</b>	<b>vD,vA,vB</b>	OR the contents of <b>vA</b> with <b>vB</b> and place the result into <b>vD</b> .



**Table 4-4. Vector Integer Logical Instructions (Continued)**

Name	Mnemonic	Syntax	Operation
Vector Logical XOR	<b>vxor</b>	<b>vD,vA,vB</b>	XOR the contents of <b>vA</b> with <b>vB</b> and place the result into <b>vD</b> .
Vector Logical AND with Complement	<b>vandc</b>	<b>vD,vA,vB</b>	AND the contents of <b>vA</b> with the complement of <b>vB</b> and place the result into <b>vD</b> .
Vector Logical NOR	<b>vnor</b>	<b>vD,vA,vB</b>	NOR the contents of <b>vA</b> with <b>vB</b> and place the result into <b>vD</b> .

### 4.2.1.5 Vector Integer Rotate and Shift Instructions

The vector integer rotate instructions are summarized in Table 4-5.

**Table 4-5. Vector Integer Rotate Instructions**

Name	Mnemonic	Syntax	Operation
Vector Rotate Left Integer	<b>vrlb</b> <b>vrlh</b> <b>vrlw</b>	<b>vD,vA,vB</b>	<p>Rotate each element in <b>vA</b> left by the number of bits specified in the low-order <math>\log_2(n)</math> bits of the corresponding element in <b>vB</b>. Place the result into the corresponding element of <b>vD</b>.</p> <p>For <b>b</b>, byte, integer length = 8 bits = 1 byte, use 16 integers from <b>vA</b> with 16 integers from <b>vB</b></p> <p>For <b>h</b>, half word, integer length = 16 bits = 2 bytes, use 8 integers from <b>vA</b> with 8 integers from <b>vB</b></p> <p>For <b>w</b>, word, integer length = 32 bits = 4 bytes, use 4 integers from <b>vA</b> with 4 integers from <b>vB</b></p>

The vector integer shift instructions are summarized in Table 4-6

**Table 4-6. Vector Integer Shift Instructions**

Name	Mnemonic	Syntax	Operation
Vector Shift Left Integer	<b>vsib</b> <b>vsih</b> <b>vsilw</b>	<b>vD,vA,vB</b>	<p>Shift each element in <b>vA</b> left by the number of bits specified in the low-order <math>\log_2(n)</math> bits of the corresponding element in <b>vB</b>. If bits are shifted out of bit 0 of the element they are lost. Supply zeros to the vacated bits on the right. Place the result into the corresponding element of <b>vD</b></p> <p>For <b>b</b>, byte, integer length = 8 bits = 1 byte, use 16 integers from <b>vA</b> with 16 integers from <b>vB</b></p> <p>For <b>h</b>, half word, integer length = 16 bits = 2 bytes, use 8 integers from <b>vA</b> with 8 integers from <b>vB</b></p> <p>For <b>w</b>, word, integer length = 32 bits = 4 bytes, use 4 integers from <b>vA</b> with 4 integers from <b>vB</b></p>

Name	Mnemonic	Syntax	Operation
Vector Shift Right Integer	<b>vsrb</b> <b>vsrh</b> <b>vsrw</b>	<b>vD,vA,vB</b>	<p>Shift each element in <b>vA</b> right by the number of bits specified in the low-order <math>\log_2(n)</math> bits of the corresponding element in <b>vB</b>. If bits are shifted out of bit <math>n-1</math> of the element they are lost. Supply zeros to the vacated bits on the left. Place the result into the corresponding element of <b>vD</b>.</p> <p>For <b>b</b>, byte, integer length = 8 bits = 1 byte, use 16 integers from <b>vA</b> with 16 integers from <b>vB</b></p> <p>For <b>h</b>, half word, integer length = 16 bits = 2 bytes, use 8 integers from <b>vA</b> with 8 integers from <b>vB</b></p> <p>For <b>w</b>, word, integer length = 32 bits = 4 bytes, use 4 integers from <b>vA</b> with 4 integers from <b>vB</b></p>
Vector Shift Right Algebraic Integer	<b>vsrab</b> <b>vsrah</b> <b>vsraw</b>	<b>vD,vA,vB</b>	<p>Shift each element in <b>vA</b> right by the number of bits specified in the low-order <math>\log_2(n)</math> bits of the corresponding element in <b>vB</b>. If bits are shifted out of bit <math>n-1</math> of the element they are lost. Replicate bit 0 of the element to fill the vacated bits on the left. Place the result into the corresponding element of <b>vD</b>.</p> <p>For <b>b</b>, byte, integer length = 8 bits = 1 byte, use 16 integers from <b>vA</b> with 16 integers from <b>vB</b></p> <p>For <b>h</b>, half word, integer length = 16 bits = 2 bytes, use 8 integers from <b>vA</b> with 8 integers from <b>vB</b></p> <p>For <b>w</b>, word, integer length = 32 bits = 4 bytes, use 4 integers from <b>vA</b> with 4 integers from <b>vB</b></p>

## 4.2.2 Vector Floating-Point Instructions

This section describes the vector floating-point instructions, that include the following:

- Vector floating-point arithmetic instructions
- Vector floating-point rounding and conversion instructions
- Vector floating-point compare instructions
- Vector floating-point estimate instructions

The AltiVec floating-point data format complies with the ANSI/IEEE-754 standard. A quantity in this format represents: a signed normalized number, a signed denormalized number, a signed zero, a signed infinity, a quiet not a number (QNaN), or a signalling NaN (SNaN). Operations perform to a Java/IEEE/C9X-compliant subset of the IEEE standard, for further details on the Java or Non-Java mode see Section 3.2.1, “Floating-Point Modes.” The AltiVec ISA does not report IEEE exceptions but rather produces default results as specified by the Java/IEEE/C9X Standard, for further details on exceptions see Section 3.2.4, “Floating-Point Exceptions.”

### 4.2.2.1 Floating-Point Division and Square-Root

AltiVec instructions do not have division or square-root instructions. AltiVec ISA implements Vector Reciprocal Estimate Floating-Point (**vrefp**) and Vector Reciprocal-Square-Root Estimate Floating-Point (**vsqrtefp**) instructions along with a Vector Negative Multiply-Subtract Floating-Point (**vnmsubfp**) instruction assisting in the Newton-Raphson refinement of the estimates. To accomplish division simply multiply the dividend ( $x/y = x * 1/y$ ) and square-root by multiplying the original number ( $\sqrt{x} = x * 1/\sqrt{x}$ ). In this way, the

AltiVec ISA provides inexpensive divides and square-roots that are fully pipelined, sub-operation scheduled, and faster even than many hardware dividers. Methods are available to further refine these to correct IEEE results, where necessary at the cost of additional software overhead.

#### 4.2.2.1.1 Floating-Point Division

The Newton-Raphson refinement step for the reciprocal  $1/B$  looks like this:

$$y1 = y0 + y0*(1 - B*y0), \text{ where } y0 = \text{recip\_est}(B)$$

This is implemented in the AltiVec ISA as follows:

```
y0 = vrefp(B)
t = vnmsubfp(y0,B,1)
y1 = vmaddfp(y0,t,y0)
```

This produces a result accurate to almost 24 bits of precision (except in the case where B is a sufficiently small denormalized number that **vrefp** generates an infinity, that, if important, must be explicitly guarded against).

To get a correctly rounded IEEE quotient from the above result, a second Newton-Raphson iteration is performed to get a correctly rounded reciprocal (y2) to the required 24 bits of precision, then the residual:

$$R = A - B*Q$$

is computed with **vnmsubfp** (where A is the dividend, B the divisor, and Q an approximation of the quotient from  $A*y2$ ). The correctly rounded quotient can then be obtained from:

$$Q' = Q + R*y2$$

The additional accuracy provided by the fused nature of the AltiVec instruction multiply-add is essential to producing the correctly rounded quotient by this method.

The second Newton-Raphson iteration may ultimately not be needed but more work must be done to show that the absolute error after the first refinement step would always be less than 1 ulp, that is a requirement of this method.

#### 4.2.2.1.2 Floating-Point Square-Root

The Newton-Raphson refinement step for reciprocal square root looks like:

$$y1 = y0 + 0.5*y0*(1 - B*y0*y0), \text{ where } y0 = \text{recip\_sqrt\_est}(B)$$

That can be implemented as follows:

```
y0 = vrsqrtefp(B)
t0 = vmaddfp(y0,y0,0.0)
t1 = vmaddfp(y0,0.5,0.0)
t0 = vnmsubfp(B,t0,1)
y1 = vmaddfp(t0,t1,y0)
```

Various methods can further refine a correctly rounded IEEE result—all more elaborate than the simple residual correction for division, and therefore are not presented here, but most of which also benefit from the negative multiply-subtract instruction.

#### 4.2.2.2 Floating-Point Arithmetic Instructions

The floating-point arithmetic instructions are summarized in Table 4-7.

**Table 4-7. Floating-Point Arithmetic Instructions**

Name	Mnemonic	Syntax	Operation
Vector Add Floating-Point	<b>vaddfp</b>	<b>vD,vA,vB</b>	Add the 4-word (32-bit) floating-point elements in <b>vA</b> to the 4-word (32-bit) floating-point elements in <b>vB</b> . Round the four intermediate results to the nearest single-precision number and placed into <b>vD</b> .
Vector Subtract Floating-Point	<b>vsubfp</b>	<b>vD,vA,vB</b>	The 4-word (32-bit) floating-point values in <b>vB</b> are subtracted from the 4 32-bit values in <b>vA</b> . The four intermediate results are rounded to the nearest single-precision floating-point and placed into <b>vD</b> .
Vector Maximum Floating-Point	<b>vmaxfp</b>	<b>vD,vA,vB</b>	Compare each of the 4 single-precision word elements in <b>vA</b> to the corresponding 4 single-precision word elements in <b>vB</b> For each of the four elements, place the larger value within each pair into <b>vD</b> . <b>vmaxfp</b> is sensitive to the sign of 0.0. When both operands are $\pm 0.0$ : $\max(+0.0, \pm 0.0) = \max(\pm 0.0, +0.0) \Rightarrow +0.0$ $\max(-0.0, -0.0) \Rightarrow -0.0$ $\max(\text{NaN}, x) \Rightarrow \text{QNaN}$ where $x = \text{any value}$
Vector Minimum Floating-Point	<b>vminfp</b>	<b>vD,vA,vB</b>	Compare each of the 4 single-precision word elements in <b>vA</b> to the corresponding 4 single-precision word elements in <b>vB</b> For each of the four elements, place the smaller value within each pair into <b>vD</b> . <b>vminfp</b> is sensitive to the sign of 0.0. When both operands are $\pm 0.0$ : $\min(-0.0, \pm 0.0) = \min(\pm 0.0, -0.0) \Rightarrow -0.0$ $\min(+0.0, +0.0) \Rightarrow +0.0$ $\min(\text{NaN}, x) \Rightarrow \text{QNaN}$ where $x = \text{any value}$

#### 4.2.2.3 Floating-Point Multiply-Add Instructions

Vector multiply-add is critically important to performance since a multiply followed by a data dependent addition is the most common idiom in DSP algorithms. In most implementations, floating-point multiply-add will perform with the same latency as either a multiply or add alone, thus doubling performance in comparing to the otherwise serial multiply and adds.

Altivec floating-point multiply-adds fuse (a multiply-add fuse implies that the full product participates in the add operation without rounding, only the final result rounds). This not only simplifies the implementation and reduces latency (by eliminating the intermediate rounding) but also increases the accuracy compared to separate multiply and adds.

Be careful as Java-compliant programs can not use multiply-add fused directly because Java requires both the product and sum to round separately. Thus to achieve strict Java compliance, perform the multiply and add with separate instructions.

To realize multiply in the AltiVec ISA use multiply-add with a zero addend (for example, **vmaddfp vD,vA,vC,vB** where (**vB** = 0.0)).

Note that in order to use multiply-add to perform an IEEE or Java-compliant multiply, the addend must be -0.0. This is necessary to insure that the sign of a zero result is correct when the product is either +0.0 or -0.0 ( $+0.0 + -0.0 \Rightarrow +0.0$ , and  $-0.0 + -0.0 \Rightarrow -0.0$ ). When the sign of a resulting 0.0 is not important, then use +0.0 as the addend that may, in some cases, avoiding the need for a second register to hold a -0.0 in addition to the integer 0/floating-point +0.0 that may already be available.

The floating-point multiply-add instructions are summarized in Table 4-8.

**Table 4-8. Floating-Point Multiply-Add Instructions**

Name	Mnemonic	Syntax	Operation
Vector Multiply-Add Floating-Point	<b>vmaddfp</b>	<b>vD,vA,vC,vB</b>	Multiply the four word floating-point elements in <b>vA</b> by the corresponding four word elements in <b>vC</b> . Add the four word elements in <b>vB</b> to the four intermediate products. Round the results to the nearest single-precision numbers and place the corresponding word elements into <b>vD</b> .
Vector Negative Multiply-Subtract Floating-Point	<b>vmmsubfp</b>	<b>vD,vA,vC,vB</b>	Multiply the four word floating-point elements in <b>vA</b> by the corresponding four word elements in <b>vC</b> . Subtract the four word floating-point elements in <b>vB</b> from the four intermediate products and invert the sign of the difference. Round the results to the nearest single-precision numbers and place the corresponding word elements into <b>vD</b> .

#### 4.2.2.4 Floating-Point Rounding and Conversion Instructions

All AltiVec floating-point arithmetic instructions use the IEEE default rounding mode, round-to-nearest. The AltiVec ISA does not provide the IEEE directed rounding modes.

The AltiVec ISA provides separate instructions for converting floating-point numbers to integral floating-point values for all IEEE rounding modes as follows:

- Round-to-nearest (**vrfin**) (round),
- Round-toward-zero (**vrfiz**) (truncate),
- Round-toward-minus-infinity (**vrfim**) (floor)
- Round-toward-positive-infinity (**vrfig**) (ceiling).

Floating-point conversions to integers (**vctuxs**, **vctxsx**) use round-toward-zero (truncate).

The floating-point rounding instructions are shown in Table 4-9.

**Table 4-9. Floating-Point Rounding and Conversion Instructions**

Name	Mnemonic	Syntax	Operation
Vector Round to Floating-Point Integer Nearest	<b>fvrfin</b>	<b>vD,vB</b>	Round to the nearest the four word floating-point elements in <b>vB</b> and place the four corresponding word elements into <b>vD</b> .
Vector Round to Floating-Point Integer toward Zero	<b>fvrfiz</b>	<b>vD,vB</b>	Round towards zero the four word floating-point elements in <b>vB</b> and place the four corresponding word elements into <b>vD</b> .
Vector Round to Floating-Point Integer toward Positive Infinity	<b>fvrfig</b>	<b>vD,vB</b>	Round towards +Infinity the four word floating-point elements in <b>vB</b> and place the four corresponding word elements into <b>vD</b> .
Vector Round to Floating-Point Integer toward Minus Infinity	<b>fvrfim</b>	<b>vD,vB</b>	Round towards -Infinity the four word floating-point elements in <b>vB</b> and place the four corresponding word elements into <b>vD</b> .
Vector Convert from Unsigned Fixed-Point Word	<b>vcfux</b>	<b>vD,vB, UIMM</b>	Convert each of the four unsigned fixed-point integer word elements in <b>vB</b> to the nearest single-precision value. Divide the result by $2^{UIMM}$ and place into the corresponding word element of <b>vD</b> .
Vector Convert from Signed Fixed-Point Word	<b>vcfsx</b>	<b>vD,vB, UIMM</b>	Convert each signed fixed-point integer word element in <b>vB</b> to the nearest single-precision value. Divide the result by $2^{UIMM}$ and place into the corresponding word element of <b>vD</b> .
Vector Convert to Unsigned Fixed-Point Word Saturate	<b>vctuxs</b>	<b>vD,vB, UIMM</b>	Multiply each of the four single-precision word elements in <b>vB</b> by $2^{UIMM}$ . The products are converted to unsigned fixed-point integers using the Round toward Zero mode. If the intermediate results are $> 2^{32}-1$ saturate to $2^{32}-1$ and if it is $< 0$ saturate to 0. Place the unsigned integer results into the corresponding word elements of <b>vD</b> .
Vector Convert to Signed Fixed-Point Word Saturate	<b>vctxsx</b>	<b>vD,vB, UIMM</b>	Multiply each of the four single-precision word elements in <b>vB</b> by $2^{UIMM}$ . The products are converted to signed fixed-point integers using Round toward Zero mode. If the intermediate results are $> 2^{32}-1$ saturate to $2^{32}-1$ and if it is $< -2^{31}$ saturate to $-2^{31}$ . Place the signed integer results into the corresponding word elements of <b>vD</b> .

### 4.2.2.5 Floating-Point Compare Instructions

The following sections describe floating-point compare instructions.

#### 4.2.2.5.1 Unordered Compares

All AltiVec floating-point compare instructions (**vcmpeqfp**, **vcmpgtfp**, **vcmpgefp**, and **vcmpbfp**) return FALSE if either operand is a NaN. Not equal-to, Not greater-than, Not greater-than-or-equal-to, and Not-in-bounds NaNs compare to everything, including themselves.

Compares always return a Boolean mask (TRUE = 0xFFFF\_FFFF, FALSE = 0x\_0000\_0000) and never return a NaN. The **vcmpeqfp** instruction is recommended as the Isnan(**vX**) test. No explicit unordered compare instructions or traps are provided. However,

the greater-than-or-equal-to predicate ( $\geq$ ) (**vcmpgefp**) is provided—in addition to the  $>$  and  $=$  predicates available for integer comparison—specifically to enable IEEE unordered comparison that would not be possible with just the  $>$  and  $=$  predicates. Table 4-10 lists the six common mathematical predicates and how they would be realized in AltiVec code.

**Table 4-10. Common Mathematical Predicates**

Case	Mathematical Predicate	AltiVec Realization	Relations			
			a>b	a<b	a=b	?
1	$a = b$	$a = b$	F	F	T	F
2	$a \neq b$ ( $?<>$ )	$\neg (a = b)$	T	T	F	T
3	$a > b$	$a > b$	T	F	F	F
4	$a < b$	$b > a$	F	T	F	F
5	$a \geq b$	$\neg (b > a)$	T	F	T	*T
6	$a \leq b$	$\neg (a > b)$	F	T	T	*T
5a	$a \geq b$	$a \geq b$	T	F	T	F
6a	$a \leq b$	$b \geq a$	F	T	T	F

\* **Note:** cases 5 and 6 implemented with greater-than (**vcmpgtfp** and **vnor**) would not yield the correct IEEE result when the relation is unordered.

Table 4-11 shows the remaining eight useful predicates and how they might be realized in AltiVec code.

**Table 4-11. Other Useful Predicates**

Case	Predicate	AltiVec Realization	Relations			
			a>b	a<b	a=b	?
7	$a ? b$	$\neg ((a=b) \vee (b>a) \vee (a>b))$	F	F	F	T
8	$a <> b$	$(a \geq b) \oplus (b \geq a)$	T	T	F	F
9	$a <=> b$	$(a \geq b) \vee (b \geq a)$	T	T	T	F
10	$a ?> b$	$\neg (b \geq a)$	T	F	F	T
11	$a ?>= b$	$\neg (b > a)$	T	F	T	T
12	$a ?< b$	$\neg (a \geq b)$	F	T	F	T
13	$a ?<= b$	$\neg (a > b)$	F	T	T	T
14	$a ?= b$	$\neg ((a > b) \vee (b > a))$	F	F	T	T

The vector floating-point compare instructions compares the elements in two vector registers word-by-word, interpreting the elements as single-precision numbers. With the exception of the Vector Compare Bounds Floating-Point (**vcmpbfp**) instruction they set the



target vector register, and CR[6] if Rc = 1, in the same manner as do the vector integer compare instructions.

The Vector Compare Bounds Floating-Point (**vcmpbfp**) instruction sets the target vector register, and CR[6] if Rc = 1, to indicate whether the elements in **vA** are within the bounds specified by the corresponding element in **vB**, as explained in the instruction description. A single-precision value *x* is said to be within the bounds specified by a single-precision value *y* if  $(-y \leq x \leq y)$ .

The floating-point compare instructions are summarized in Table 4-12.

**Table 4-12. Floating-Point Compare Instructions**

Name	Mnemonic	Syntax	Operation
Vector Compare Greater Than Floating-Point [Record]	<b>vcmpgtfp</b> [.]	<b>CR6, vD, vA, vB</b>	Compare each of the 4 single-precision word elements in <b>vA</b> to the corresponding four single-precision word elements in <b>vB</b> For each element, if <b>vA</b> > <b>vB</b> then set the corresponding element in <b>vD</b> to all 1's otherwise clear the element in <b>vD</b> to all 0's If the record bit (Rc = 1) is set in the vector compare instruction, then <b>vD</b> == 1, (all elements true) then CR6[0] is set <b>vD</b> == 0, (all elements false) then CR6[2] is set
Vector Compare Equal to Floating-Point [Record]	<b>vcmpeqfp</b> [.]	<b>CR6, vD, vA, vB</b>	Compare each of the 4 single-precision word elements in <b>vA</b> to the corresponding 4 single-precision word elements in <b>vB</b> . For each element, if <b>vA</b> = <b>vB</b> then set the corresponding element in <b>vD</b> to all 1's otherwise clear the element in <b>vD</b> to all 0's If the record bit (Rc = 1) is set in the vector compare instruction then <b>vD</b> == 1, (all elements true) then CR6[0] is set <b>vD</b> == 0, (all elements false) then CR6[2] is set
Vector Compare Greater Than or Equal to Floating-Point [Record]	<b>vcmpgeqfp</b> [.]	<b>CR6, vD, vA, vB</b>	Compare each of the 4 single-precision word elements in <b>vA</b> to the corresponding 4 single-precision word elements in <b>vB</b> . For each element, if <b>vA</b> >= <b>vB</b> then set the corresponding element in <b>vD</b> to all 1's otherwise clear the element in <b>vD</b> to all 0's If the record bit (Rc = 1) is set in the vector compare instruction then <b>vD</b> == 1, (all elements true) then CR6[0] is set <b>vD</b> == 0, (all elements false) then CR6[2] is set



**Table 4-12. Floating-Point Compare Instructions**

Name	Mnemonic	Syntax	Operation
Vector Compare Bounds Floating-Point [Record]	<b>vcmpbfp</b> [.]	<b>CR6,vD,vA,vB</b>	<p>Compare each of the 4 single-precision word elements in <b>vA</b> to the corresponding single-precision word elements in <b>vB</b>. A 2-bit value is formed that indicates whether the element in <b>vA</b> is within the bounds specified by the element in <b>vB</b>, as follows.</p> <p>Bit 0 of the two-bit value is cleared to 0 if the element in <b>vA</b> is <math>\leq</math> to the element in <b>vB</b>, and is set otherwise.</p> <p>Bit 1 of the two-bit value is cleared to 0 if the element in <b>vA</b> is <math>\geq</math> to the negation of the element in <b>vB</b>, and is set otherwise.</p> <p>The two-bit value is placed into the high-order two bits of the corresponding word element of <b>vD</b> and the remaining bits of the element are cleared to 0.</p> <p>If <math>Rc=1</math>, <b>CR6[2]</b> is set to 1 when all four elements in <b>vA</b> are within the bounds specified by the corresponding element in <b>vB</b></p>

#### 4.2.2.6 Floating-Point Estimate Instructions

The floating-point estimate instructions are summarized in Table 4-13.

**Table 4-13. Floating-Point Estimate Instructions**

Name	Mnemonic	Syntax	Operation
Vector Reciprocal Estimate Floating-Point	<b>vrefp</b>	<b>vD,vB</b>	Place estimates of the reciprocal of each of the four word floating-point source elements in <b>vB</b> in the corresponding four word elements in <b>vD</b> .
Vector Reciprocal Square Root Estimate Floating-Point	<b>vrqrtefp</b>	<b>vD,vB</b>	Place estimates of the reciprocal square-root of each of the four word source elements in <b>vB</b> in the corresponding four word elements in <b>vD</b> .
Vector Log2 Estimate Floating-Point	<b>vlogefp</b>	<b>vD,vB</b>	Place estimates of the base 2 logarithm of each of the four word source elements in <b>vB</b> in the corresponding four word elements in <b>vD</b> .
Vector 2 Raised to the Exponent Estimate Floating-Point	<b>vexpteftp</b>	<b>vD,vB</b>	Place estimates of 2 raised to the power of each of the four word source elements in <b>vB</b> in the corresponding four word elements in <b>vD</b> .

#### 4.2.3 Load and Store Instructions

Only very basic load and store operations are provided in the AltiVec ISA. This keeps the circuitry in the memory path fast so the latency of memory operations will be low. Instead, a powerful set of field manipulation instructions are provided to manipulate data into the desired alignment and arrangement after the data has been brought into the vector registers.

Load vector indexed (**lvx**, **lvxl**) and store vector indexed (**stvx**, **stvxl**) instructions transfer an aligned quad-word vector between memory and vector registers. Load vector element indexed (**lvebx**, **lvehx**, **lvewx**) and store vector element indexed instructions (**stvebx**,

**stvehx**, **stvewx**) transfer byte, half-word, and word scalar elements between memory and vector registers.

All vector loads and vector stores use the index ( $\mathbf{rA}|0 + \mathbf{rB}$ ) addressing mode to specify the target memory address. The AltiVec ISA does not provide any update forms. A **lvehx**, **lvehx**, or **lvewx** instruction transfers a scalar data element from memory into the destination vector register, leaving other elements in the vector with boundedly-undefined values. A **stvehx**, **stvehx**, or **stvewx** instruction transfers a scalar data element from the source vector register to memory leaving other elements in the quad word unchanged. No data alignment occurs, that is, all scalar data elements are transferred directly on their natural memory byte-lanes to or from the corresponding element in the vector register. Quad word memory accesses made by **lvx**, **lvxl**, **stvx**, and **stvxl** instructions are not guaranteed to be atomic. Direct-store segments ( $T=1$ ) are not supported. Any vector load or store that attempts to access a direct-store segment will cause a data storage interrupt (DSI) exception.

#### 4.2.3.1 Alignment

All memory references must be size aligned. If a vector load or store address is not properly size aligned, the suitable number of least significant bits are ignored, and a size aligned transfer occurs instead. Data alignment must be performed explicitly after being brought into the registers. No assistance is provided to assist in aligning individual scalar elements that are not aligned on their natural size boundary. However, assistance is provided for justifying non-size-aligned vectors. This is provided through the special Load Vector for Shift Left (**lvsl**) and Load Vector for Shift Right (**lvslr**) instructions that compute the proper Vector Permute (**vperm**) control vector from the misaligned memory address. For details on how to use these instructions to align data see Section 3.1.6, “Quad-Word Data Alignment.”

The **lvx**, **lvxl**, **stvx**, and **stvxl** instructions can be used to move all sorts of data, not just multimedia data, in typical PowerPC environments. Therefore, because vector loads and stores are size-aligned, care should be taken to align data on even quad-word boundaries for maximum performance.

#### 4.2.3.2 Load and Store Address Generation

Vector load and store operations generate effective addresses using register indirect with index mode.

##### 4.2.3.2.1 Register Indirect with Index Addressing for Loads and Stores

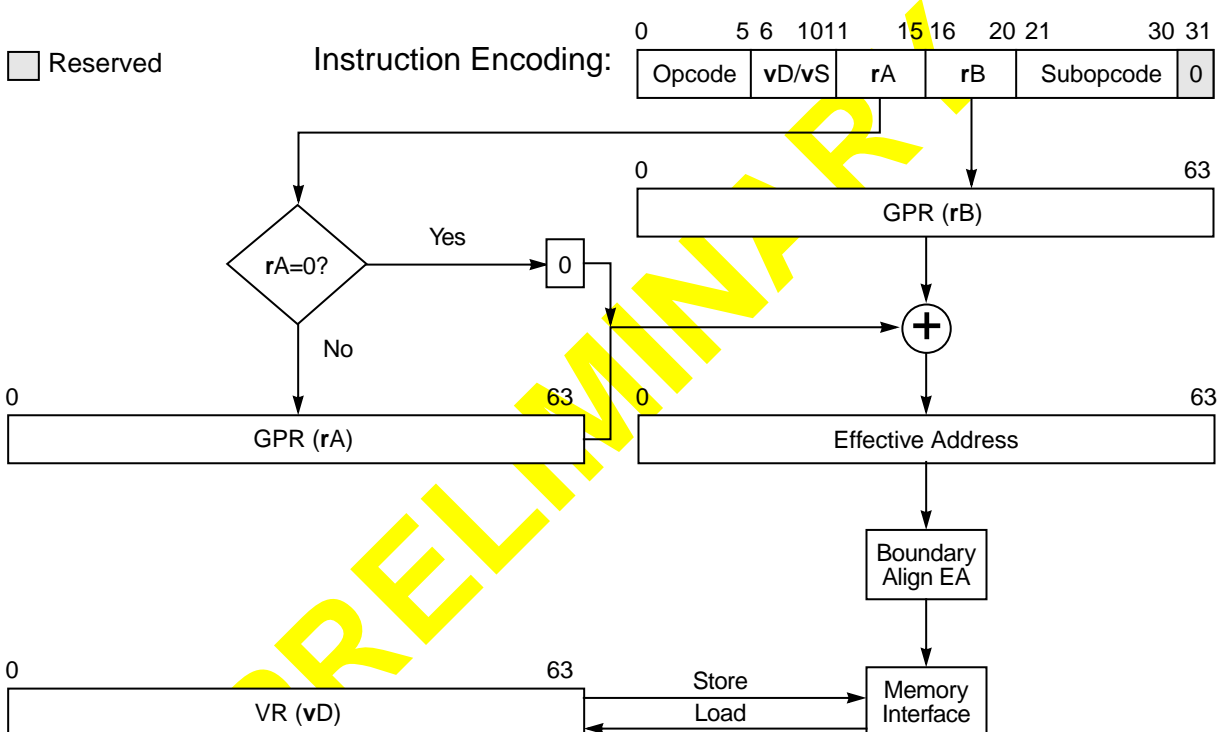
All AltiVec vx load and store instructions use register indirect with index addressing mode that cause the contents of two general-purpose registers (specified as operands **rA** and **rB**) to be added in the generation of the effective address (EA). A zero in place of the **rA** operand causes a zero to be added to the contents of the GPR specified in **rB**. The option to specify **rA** or 0 is shown in the instruction descriptions as ( $\mathbf{rA}|0$ ). If when combining addresses ( $\mathbf{rA}|0 + \mathbf{rB}$ ) the address becomes unaligned, for a half word, word, or quad word,

the effective address is ANDed with the appropriate zero values to boundary align the address and is summarized in Table 4-14.

**Table 4-14. Effective Address Alignment**

Operand	Effective Address Bit	Setting
Indexed Half word	EA[63]	0b0
Indexed Word	EA[62–63]	0b00
Indexed Quad word	EA[60–63]	0b0000

Figure 4-1 shows how an effective address is generated when using register indirect with index addressing.



**Figure 4-1. Register Indirect with Index Addressing for Loads/Stores**

#### 4.2.3.3 Vector Load Instructions

For vector load instructions, the byte, half word, or word addressed by the EA (effective address) is loaded into **rD**.

The default byte and bit ordering is big-endian as in the PowerPC architecture; see Section 3.1.2, “Byte Ordering,” for information about little-endian byte ordering.

Table 4-15 summarizes the vector load instructions.

**Table 4-15. Integer Load Instructions**

Name	Mnemonic	Syntax	Operation
Load Vector Element Integer Indexed	<b>lvebx</b> <b>lvehx</b> <b>lviewx</b>	<b>vD,rA,rB</b>	<p>The EA is the sum <math>(rA[0] + (rB))</math>. Load the byte, half word, or word in memory addressed by the EA into the low-order bits of <b>vD</b>. The remaining bits in <b>vD</b> are set to boundedly undefined values.</p> <p>Because memory must stay aligned, the EA is set to default to alignment:</p> <p>For <b>b</b>, byte, integer length = 8 bits = 1 byte,</p> <p>For <b>h</b>, half word, integer length = 16 bits = 2 bytes, EA[62-63] is set to 0b0</p> <p>For <b>w</b>, word, integer length= 32 bits = 4 bytes, EA[61-63] is set to 0b00</p>
Load Vector Element Indexed	<b>lvx</b>	<b>vD,rA,rB</b>	<p>The EA is the sum <math>(rA[0] + (rB))</math>. Load the double word in memory addressed by the EA into <b>vD</b>.</p> <p>Because memory needs to stay aligned, the EA is set to default to alignment:</p> <p>For <b>q</b>, quad word, integer length =128 bits = 8 bytes, the EA[60-63] is set to 0b0000</p> <p>LRU = 0</p> <p>If the processor is in little-endian mode, load the double word in memory addressed by EA into <b>vD</b>[6–127] and load the double word in memory addressed by EA+8 into <b>vD</b>[0–63].</p>
Load Vector Element Indexed LRU	<b>lvxl</b>	<b>vD,rA,rB</b>	<p>The EA is the sum <math>(rA[0] + (rB))</math>. Load the double word in memory addressed by the EA into <b>vD</b>.</p> <p>For <b>d</b>, double word, integer length=64 bits = 4 bytes, the EA[60-63] is set to 0b0000</p> <p>LRU =1, least recently used, hints that the quad word in the memory addressed by EA will probably not be needed again by the program in the near future.</p> <p>If the processor is in little-endian mode, load the double word in memory addressed by EA into <b>vD</b>[64–127] and load the double word in memory addressed by EA+8 into <b>vD</b>[0–63].</p>

#### 4.2.3.3.1 Vector Load Instructions Supporting Alignment

The **lvsl** and **lvslr** instructions can be used to create the permute control vector to be used by a subsequent **vperm** instruction. Let X and Y be the contents of **vA** and **vB** specified by **vperm**. The control vector created by **lvsl** causes the **vperm** to select the high-order 16 bytes of the result of shifting the 32-byte value  $X \parallel Y$  left by sh bytes (sh = the value in EA[60-63]). The control vector created by **lvslr** causes the **vperm** to select the low-order 16 bytes of the result of shifting  $X \parallel Y$  right by sh bytes.

These instructions can also be used to rotate or shift the contents of a vector register left **lvsl** or right **lvslr** by sh bytes. For rotating, the vector register to be rotated should be specified as both the **vA** and the **vB** register for **vperm**. For shifting left, the **VB** register for **vperm** should be a register containing all zeros and **vA** should contain the value to be shifted, and vice versa for shifting right. For further examples on how to align the data see Section 3.1.6, “Quad-Word Data Alignment.” The default byte and bit ordering is big-endian as in the

PowerPC architecture; see Section 3.1.2.2, “Little-Endian Byte Ordering,” for information about little-endian byte ordering.

Table 4-15 summarizes the vector alignment instructions.

**Table 4-16. Vector Load Instructions Supporting Alignment**

Name	Mnemonic	Syntax	Operation
Load Vector for Shift Left	<b>lvsl</b>	<b>vD,rA,rB</b>	<p>The EA is the sum (rA[0] + (rB)). The EA[60–63] = sh, then based on a table lookup place the value in vD</p> <p>if sh = 0x0 then (vD)0:127 &lt;- 0x000102030405060708090A0B0C0D0E0F  if sh = 0x1 then (vD)0:127 &lt;- 0x0102030405060708090A0B0C0D0E0F10  if sh = 0x2 then (vD)0:127 &lt;- 0x02030405060708090A0B0C0D0E0F1011  if sh = 0x3 then (vD)0:127 &lt;- 0x030405060708090A0B0C0D0E0F101112  if sh = 0x4 then (vD)0:127 &lt;- 0x0405060708090A0B0C0D0E0F10111213  if sh = 0x5 then (vD)0:127 &lt;- 0x05060708090A0B0C0D0E0F1011121314  if sh = 0x6 then (vD)0:127 &lt;- 0x060708090A0B0C0D0E0F101112131415  if sh = 0x7 then (vD)0:127 &lt;- 0x0708090A0B0C0D0E0F10111213141516  if sh = 0x8 then (vD)0:127 &lt;- 0x08090A0B0C0D0E0F1011121314151617  if sh = 0x9 then (vD)0:127 &lt;- 0x090A0B0C0D0E0F101112131415161718  if sh = 0xA then (vD)0:127 &lt;- 0x0A0B0C0D0E0F10111213141516171819  if sh = 0xB then (vD)0:127 &lt;- 0x0B0C0D0E0F101112131415161718191A  if sh = 0xC then (vD)0:127 &lt;- 0x0C0D0E0F101112131415161718191A1B  if sh = 0xD then (vD)0:127 &lt;- 0x0D0E0F101112131415161718191A1B1C  if sh = 0xE then (vD)0:127 &lt;- 0x0E0F101112131415161718191A1B1C1D  if sh = 0xF then (vD)0:127 &lt;- 0x0F101112131415161718191A1B1C1D1E</p>
Load Vector for Shift Right	<b>lvsr</b>	<b>vD,rA,rB</b>	<p>The EA is the sum (rA[0] + (rB)). The EA[60–63] = sh, then based on the table lookup below place the value in vD</p> <p>if sh = 0x0 then (vD)0:127 &lt;- 0x101112131415161718191A1B1C1D1E1F  if sh = 0x1 then (vD)0:127 &lt;- 0x0F101112131415161718191A1B1C1D1E  if sh = 0x2 then (vD)0:127 &lt;- 0x0E0F101112131415161718191A1B1C1D  if sh = 0x3 then (vD)0:127 &lt;- 0x0D0E0F101112131415161718191A1B1C  if sh = 0x4 then (vD)0:127 &lt;- 0x0C0D0E0F101112131415161718191A1B  if sh = 0x5 then (vD)0:127 &lt;- 0x0B0C0D0E0F101112131415161718191A  if sh = 0x6 then (vD)0:127 &lt;- 0x0A0B0C0D0E0F10111213141516171819  if sh = 0x7 then (vD)0:127 &lt;- 0x090A0B0C0D0E0F101112131415161718  if sh = 0x8 then (vD)0:127 &lt;- 0x08090A0B0C0D0E0F1011121314151617  if sh = 0x9 then (vD)0:127 &lt;- 0x0708090A0B0C0D0E0F10111213141516  if sh = 0xA then (vD)0:127 &lt;- 0x060708090A0B0C0D0E0F101112131415  if sh = 0xB then (vD)0:127 &lt;- 0x05060708090A0B0C0D0E0F1011121314  if sh = 0xC then (vD)0:127 &lt;- 0x0405060708090A0B0C0D0E0F10111213  if sh = 0xD then (vD)0:127 &lt;- 0x030405060708090A0B0C0D0E0F101112  if sh = 0xE then (vD)0:127 &lt;- 0x02030405060708090A0B0C0D0E0F1011  if sh = 0xF then (vD)0:127 &lt;- 0x0102030405060708090A0B0C0D0E0F10</p>

#### 4.2.3.4 Vector Store Instructions

For vector store instructions, the contents of vector register used as a source (**vS**) are stored into the byte, half word, word or quad word in memory addressed by the effective address (EA). Table 4-17 provides a summary of the vector store instructions.

**Table 4-17. Integer Store Instructions**

Name	Mnemonic	Syntax	Operation
Store Vector Element Integer Indexed	<b>svetbx</b> <b>svethx</b> <b>svetwx</b>	<b>vS,rA,rB</b>	The EA is the sum ( <b>rA</b> [0] + <b>rB</b> ). Store the contents of the low-order bits of <b>vS</b> into the integer in memory addressed by the EA.  Because memory needs to stay aligned, the EA is set to default to alignment: For <b>b</b> , byte, integer length = 8 bits = 1 byte, For <b>h</b> , half word, integer length = 16 bits = 2 bytes, EA[62–63] is set to 0b0 For <b>w</b> , word, integer length = 32 bits = 4 bytes, EA[61–63] is set to 0b00
Store Vector Element Indexed	<b>stvx</b>	<b>vS,rA,rB</b>	The EA is the sum ( <b>rA</b> [0] + <b>rB</b> ). Store the contents of <b>vS</b> into the quad word in memory addressed by the EA.  For <b>q</b> , quad word, integer length = 64 bits = 4 bytes, the EA[60–63] is set to 0b0000  LRU = 0  If the processor is in little-endian mode, store the contents of <b>vS</b> [64–127] into the double word in memory addressed by EA, and store the contents of <b>vS</b> [0–63] into the double word in memory addressed by EA+8.
Store Vector Element Indexed LRU	<b>stvx</b>	<b>vD,rA,rB</b>	The EA is the sum ( <b>rA</b> [0] + <b>rB</b> ). Store the contents of <b>vS</b> into the quad word in memory addressed by the EA.  For <b>d</b> , double word, integer length=64 bits = 4 bytes, the EA[60–63] is set to 0b0000  LRU = 1, least recently used, hints that the quad word in the memory addressed by EA will probably not be needed again by the program in the near future.  If the processor is in little-endian mode, store the contents of <b>vS</b> [64–127] into the double word in memory addressed by EA, and store the contents of <b>vS</b> [0–63] into the double word in memory addressed by EA+8.

#### 4.2.4 Control Flow

AltiVec vx instructions can be freely intermixed with existing PowerPC instructions to form a complete program. AltiVec vx instructions do provide a vector compare and select mechanism to implement conditional execution as the preferred mechanism to control data flow in AltiVec vx programs. And AltiVec vx vector compare instructions can update the condition register thus providing the communication from AltiVec vx execution units to PowerPC branch instructions necessary to modify program flow based on vector data.

#### 4.2.5 Vector Permutation and Formatting Instructions

Vector pack, unpack, merge, splat, permute, and select can be used to accelerate various vector math and vector formatting. Details of the various instructions follows.

### 4.2.5.1 Vector Pack Instructions

Half-word vector pack instructions (**vpkuhum**, **vpkuhus**, **vpkshus**, **vpkshss**) truncate the sixteen half words from two concatenated source operands producing a single result of sixteen bytes (quad word) using either modulo( $2^8$ ), 8-bit signed-saturation, or 8-bit unsigned-saturation to perform the truncation. Similarly, word vector pack instructions (**vpkuwum**, **vpkuwus**, **vpkswus**, **vpksws**) truncate the eight words from two concatenated source operands producing a single result of eight half words using modulo( $2^{16}$ ), 16-bit signed-saturation, or 16-bit unsigned-saturation to perform the truncation.

One special purpose form of Vector Pack Pixel (**vpkpx**) instruction is provided that packs eight 32-bit (8/8/8/8) pixels from two concatenated source operands into a single result of eight 16-bit 1/5/5/5  $\alpha$ RGB pixels. The least significant bit of the first 8-bit element becomes the 1-bit  $\alpha$  field, and each of the three 8-bit R, G, and B fields are reduced to 5 bits by discarding the 3 lsbs.

Table 4-18 describes the vector pack instructions.

**Table 4-18. Vector Pack Instructions**

Name	Mnemonic	Syntax	Operation
Vector Pack Unsigned Integer[h,w] Unsigned Modulo	<b>vpkuhum</b> <b>vpkuwum</b>	<b>vD, vA, vB</b>	Concatenate the low order unsigned integers of <b>vA</b> and the low order unsigned integers of <b>vB</b> and place into <b>vD</b> using unsigned modulo arithmetic. <b>vA</b> is placed in the lower order double word of <b>vD</b> and <b>vB</b> is placed into the higher order double word of <b>vD</b> .  For <b>h</b> , half word, integer length = 16 bits = 2 bytes, 8 unsigned integers, in other words the 8 low order bytes of the half words from <b>vA</b> and <b>vB</b>  For <b>w</b> , word, integer length = 32 bits = 4 bytes, 4 unsigned integers, in other words the 4 low order half words of the words from <b>vA</b> and <b>vB</b>
Vector Pack Unsigned Integer[h,w] Unsigned Saturate	<b>vpkuhus</b> <b>vpkuwus</b>	<b>vD, vA, vB</b>	Concatenate the low order unsigned integers of <b>vA</b> and the low order unsigned integers of <b>vB</b> and place into <b>vD</b> using unsigned saturate clamping mode. <b>vA</b> is placed in the lower order double word of <b>vD</b> and <b>vB</b> is placed into the higher order double word of <b>vD</b> .  For <b>h</b> , half word, integer length = 16 bits = 2 bytes, 8 unsigned integers, in other words the 8 low order bytes of the half words from <b>vA</b> and <b>vB</b>  For <b>w</b> , word, integer length = 32 bits = 4 bytes, 4 unsigned integers, in other words the 4 low order words of the half words from <b>vA</b> and <b>vB</b>
Vector Pack Signed Integer[h,w] Unsigned Saturate	<b>vpkshus</b> <b>vpkswus</b>	<b>vD, vA, vB</b>	Concatenate the low order signed integers of <b>vA</b> and the low order signed integers of <b>vB</b> and place into <b>vD</b> using unsigned saturate clamping mode. <b>vA</b> is placed in the lower order double word of <b>vD</b> and <b>vB</b> is placed into the higher order double word of <b>vD</b> .  For <b>h</b> , half word, integer length = 16 bits = 2 bytes, 8 signed integers, in other words the 8 low order bytes of the half word from <b>vA</b> and <b>vB</b>  For <b>w</b> , word, integer length = 32 bits = 4 bytes, 4 signed integers, in other words the 4 low order half words of the words from <b>vA</b> and <b>vB</b>



**Table 4-18. Vector Pack Instructions (Continued)**

Name	Mnemonic	Syntax	Operation
Vector Pack Signed Integer [h,w] Unsigned Saturate	<b>vpkshss</b> <b>vpkswss</b>	<b>vD, vA, vB</b>	<p>Concatenate the low order signed integers of <b>vA</b> and the low order signed integers of <b>vB</b> are concatenated and place into <b>vD</b> using signed saturate clamping mode. <b>vA</b> is placed in the lower order double word of <b>vD</b> and <b>vB</b> is placed into the higher order double word of <b>vD</b>.</p> <p>For <b>h</b>, half word, integer length = 16 bits = 2 bytes, 8 signed integers, in other words the 8 low order bytes of the half word from <b>vA</b> and <b>vB</b></p> <p>For <b>w</b>, word, integer length = 32 bits = 4 bytes, 4 signed integers, in other words the 4 low order half words of the words from <b>vA</b> and <b>vB</b></p>
Vector Pack Pixel	<b>vpkpx</b>	<b>vD, vA, vB</b>	<p>Each word element in <b>vA</b> and <b>vB</b> is packed to 16 bits and the half word is placed into <b>vD</b>. Each word from <b>vA</b> and <b>vB</b> is packed to 16 bits in the following order:</p> <p>[bit 7 of the first byte (bit 7 of the word)]</p> <p>[bits 0–4 of the second byte (bits 8–12 of the word)]</p> <p>[bits 0–4 of the third byte (bits 16–20 of the word)]</p> <p>[bits 0–4 of the fourth byte (bits 24–28 of the word)]</p> <p><b>vA</b> half words are placed in the lower order double word of <b>vD</b> and <b>vB</b> half words are placed into the higher order double word of <b>vD</b>.</p> <p>For <b>h</b>, half word, integer length = 16 bits = 2 bytes, 8 signed integers, in other words the 8 low order bytes of the half word from <b>vA</b> and <b>vB</b></p> <p>For <b>w</b>, word, integer length = 32 bits = 4 bytes, 4 signed integers, in other words the 4 low order half words of the words from <b>vA</b> and <b>vB</b></p>

#### 4.2.5.2 Vector Unpack Instructions

Byte vector unpack instructions unpack the 8 low bytes (or 8 high bytes) of one source operand into 8 half words using sign extension to fill the MSBs. Half word vector unpack instructions unpack the 4 low half words (or 4 high half words) of one source operand into 4 words using sign extension to fill the MSBs.

A special purpose form of vector unpack is provided, the Vector Unpack Low Pixel (**vupklpx**) and the Vector Unpack High Pixel (**vupkhp**) instructions for 1/5/5/5  $\alpha$ RGB pixels. The 1/5/5/5 pixel vector unpack, unpacks the four low 1/5/5/5 pixels (or four 1/5/5/5 high pixels) into four 32-bit (8/8/8/8) pixels. The 1-bit  $\alpha$  element in each pixel is sign extended to 8 bits, and the 5-bit R, G, and B elements are each zero extended to 8 bits.



Table 4-18 describes the unpack instructions.

**Table 4-19. Vector Unpack Instructions**

Name	Mnemonic	Syntax	Operation
Vector Unpack High Signed Integer	<b>vupkhsb</b> <b>vupkhsh</b>	vD, vB	Each signed integer element in the high order double word of vB is sign extended to fill the MSBs in a signed integer and then is placed into vD.  For <b>b</b> , byte, integer length = 8 bits = 1 byte, 8 signed bytes from the high order double word of vB are unpacked and sign extended to 8 half words into vD.  For <b>h</b> , half word, integer length = 16 bits = 2 bytes, 8 signed half words from the high order double word of vB are unpacked and sign extended to 4 words into vD
Vector Unpack High Pixel	<b>vupkhpX</b>	vD, vB	Each half-word element in the high order double word of vB is unpacked to produce a 32-bit word that is then placed in the same order into vD.  A half-word element is unpacked to 32 bits by concatenating, in order, the results of the following operations. [sign-extend bit 0 of the half word to 8 bits [zero-extend bits 1–5 of the half word to 8 bits [zero-extend bits 6–10 of the half word to 8 bits [zero-extend bits 11–15 of the half word to 8 bits
Vector Unpack Low Signed Integer	<b>vupklbsb</b> <b>vupklsh</b>	vD, vB	Each signed integer element in the low order double word of vB is sign extended to fill the MSBs in a signed integer and then is placed into vD.  For <b>b</b> , byte, integer length = 8 bits = 1 byte, 8 signed bytes from the low order double word of vB are unpacked and sign extended to 8 half words in vD.  For <b>h</b> , half word, integer length = 16 bits = 2 bytes, 8 signed half words from the low order double word of vB are unpacked and sign extended into 4 words in vD
Vector Unpack Low Pixel	<b>vupklpx</b>	vD, vB	Each half-word element in the low order double word of vB is unpacked to produce a 32-bit word that is then placed in the same order into vD.  A half-word element is unpacked to 32 bits by concatenating, in order, the results of the following operations. [sign-extend bit 0 of the half word to 8 bits [zero-extend bits 1–5 of the half word to 8 bits [zero-extend bits 6–10 of the half word to 8 bits [zero-extend bits 11–15 of the half word to 8 bits

### 4.2.5.3 Vector Merge Instructions

Byte vector merge instructions interleave the 8 low bytes (or 8 high bytes) from two source operands producing a result of 16 bytes. Similarly, half-word vector merge instructions interleave the 4 low half words (or 4 high half words) of two source operands producing a result of 8 half words, and word vector merge instructions interleave the 2 low words (or 2 high words) from two source operands producing a result of 4 words. The vector merge

instruction has many uses, notable among them is a way to efficiently transpose SIMD vectors. Table 4-18 describes the merge instructions.

**Table 4-20. Vector Merge Instructions**

Name	Mnemonic	Syntax	Operation
Vector Merge High Integer	<b>vmrghb</b> <b>vmrghh</b> <b>vmrghw</b>	<b>vD, vA, vB</b>	<p>Each integer element in the high order double word of <b>vA</b> is placed into the low order integer element in <b>vD</b>. Each integer element in the high order double word of <b>vB</b> is placed into the high order integer element in <b>vD</b>.</p> <p>For <b>b</b>, byte, integer length = 8 bits = 1 byte, 8 bytes from the high order double word of <b>vA</b> are placed into the low order byte of each half word in <b>vD</b> and 8 bytes from the high order double word of <b>vB</b> are placed into the high order byte of each half word in <b>vD</b>.</p> <p>For <b>h</b>, half word, integer length = 16 bits = 2 bytes, 4 half words from the high order double word of <b>vA</b> are placed into the low order half word of each word in <b>vD</b> and 4 half words from the high order double word of <b>vB</b> are placed into the high order half word of each word in <b>vD</b>.</p> <p>For <b>w</b>, word, integer length = 32 bits = 4 bytes, 2 words from the high order double word of <b>vA</b> are placed into the low order word of each double word in <b>vD</b> and 2 words from the high order double word of <b>vB</b> are placed into the high order word of each double word in <b>vD</b>.</p>
Vector Merge Low Integer	<b>vmrglb</b> <b>vmrglh</b> <b>vmrglw</b>	<b>vD, vA, vB</b>	<p>Each integer element in the low order double word of <b>vA</b> is placed into the low order integer element in <b>vD</b>. Each integer element in the low order double word of <b>vB</b> is placed into the high order integer element in <b>vD</b>.</p> <p>For <b>b</b>, byte, integer length = 8 bits = 1 byte, 8 bytes from the low order double word of <b>vA</b> are placed into the low order byte of each half word in <b>vD</b> and 8 bytes from the low order double word of <b>vB</b> are placed into the high order byte of each half word in <b>vD</b>.</p> <p>For <b>h</b>, half word, integer length = 16 bits = 2 bytes, 4 half words from the low order double word of <b>vA</b> are placed into the low order half word of each word in <b>vD</b> and 4 half words from the low order double word of <b>vB</b> are placed into the high order half word of each word in <b>vD</b>.</p> <p>For <b>w</b>, word, integer length = 32 bits = 4 bytes, 2 words from the low order double word of <b>vA</b> are placed into the low order word of each double word in <b>vD</b> and 2 words from the low order double word of <b>vB</b> are placed into the high order word of each double word in <b>vD</b>.</p>

#### 4.2.5.4 Vector Splat Instructions

When a program needs to perform arithmetic vector, the vector splat instructions can be used in preparation for performing arithmetic for which one source vector is to consist of elements that all have the same value (e.g., multiplying all elements of a Vector Register by a constant). Vector splat instructions can be used to move data where it is required. For example to multiply all elements of a vector register by a constant, the vector splat instructions can be used to splat the scalar into the vector register. Likewise, when storing a scalar into an arbitrary memory location, it must be splatted into a vector register, and that

register specified as the source of the store. This will guarantee that the data appears in all possible positions of that scalar size for the store.

**Table 4-21. Vector Splat Instructions**

Name	Mnemonic	Syntax	Operation
Vector Splat Integer	<b>vspltb</b> <b>vsplth</b> <b>vspltw</b>	<b>vD, vB, UIMM</b>	Replicate the contents of element UIMM in <b>vB</b> and place into each element in <b>vD</b> .  For <b>b</b> , byte, integer length = 8 bits = 1 byte, each element is a byte.  For <b>h</b> , half word, integer length = 16 bits = 2 bytes, each element is a half word.  For <b>w</b> , word, integer length = 32 bits = 4 bytes, 2 words each element is a word.
Vector Splat Immediate Signed Integer	<b>vspltisb</b> <b>vspltish</b> <b>vspltisw</b>	<b>vD, SIMM</b>	Sign-extend the value of the SIMM field to the length of the element and replicate that value and place into each element in <b>vD</b> .  For <b>b</b> , byte, integer length = 8 bits = 1 byte, each element is a byte.  For <b>h</b> , half word, integer length = 16 bits = 2 bytes, each element is a half word.  For <b>w</b> , word, integer length = 32 bits = 4 bytes, 2 words each element is a word.

#### 4.2.5.5 Vector Permute Instructions

Permute instructions allow any byte in any two source vector registers to be directed to any byte in the destination vector. The fields in a third source operand specify from which field in the source operands the corresponding destination field will be taken. The Vector Permute (**vperm**) instruction is a very powerful one that provides many useful functions. For example, it provides a good way to perform table-lookups and data alignment operations. An example of how to use the command in aligning data see Section 3.1.6, “Quad-Word Data Alignment.” Table 4-18 describes the vector permute instruction.

**Table 4-22. Vector Permute Instruction**

Name	Mnemonic	Syntax	Operation
Vector Permute	<b>vperm</b>	<b>vD, vA, vB, vC</b>	<b>vC</b> specifies which bytes from <b>vA</b> and <b>vB</b> are to be copied and placed into the byte elements in <b>vD</b> .

#### 4.2.5.6 Vector Select Instruction

Data flow in the vector unit can be controlled without branching by using a vector compare and the vector select (**vsel**) instructions. In this use, the compare result vector is used directly as a mask operand to vector select instructions. The **vsel** instruction selects one field from one or the other of two source operands under control of its mask operand. Use of the TRUE/FALSE compare result vector with select in this manner produces a two instruction equivalent of conditional execution on a per-field basis. Table 4-23 describes the **vsel** instruction.

**Table 4-23. Vector Select Instruction**

Name	Mnemonic	Syntax	Operation
Vector Select	<b>vsel</b>	<b>vD,vA,vB,vC</b>	For each bit, compare the value in <b>vC</b> to the value 0b0 and if it equals 0b0 then load <b>vD</b> with <b>vA</b> 's corresponding bit value otherwise compare the value in <b>vC</b> to the value 0b1 and if it equals 0b1 then load <b>vD</b> with <b>vB</b> 's corresponding bit value.

### 4.2.5.7 Vector Shift Instructions

The vector shift instructions shift the contents of a vector register or of a pair of vector registers left or right by a specified number of bytes (**vslo**, **vsro**, **vsldoi**) or bits (**vsl**, **vsr**). Depending on the instruction, this shift count is specified either by low-order bits of a vector register or by an immediate field in the instruction. In the former case the low-order 7 bits of the shift count register give the shift count in bits ( $0 \leq \text{count} \leq 127$ ). Of these 7 bits, the high-order 4 bits give the number of complete bytes by which to shift and are used by **vslo** and **vsro**; the low-order 3 bits give the number of remaining bits by which to shift and are used by **vsl** and **vsr**.

There are two methods of specifying an inter-element shift or rotate of two source vector registers, extracting 16 bytes as the result vector. There is also a method for shifting a single source vector register left or right by any number of bits.

Table 4-24 describes the various vector shift instructions.

**Table 4-24. Vector Shift Instructions**

Name	Mnemonic	Syntax	Operation
Vector Shift Left	<b>vsl</b>	<b>vD,vA,vB</b>	Shift <b>vA</b> left by the 3 lsbs of <b>vB</b> , and place the result into <b>vD</b> If <b>vB</b> value is invalid, the default result is boundedly undefined
Vector Shift Left Double by Octet Immediate	<b>vsldoi</b>	<b>vD,vA,vB,SH</b>	Shift <b>vB</b> left by the 3 lsbs of <b>SH</b> value and then OR with <b>vA</b> , place the result into <b>vD</b> If <b>vB</b> value is invalid, the default result is 0
Vector Shift Left by Octet	<b>vslo</b>	<b>vD,vA,vB</b>	Shift <b>vA</b> left by the 3 lsbs of <b>vB</b> , and place the result into <b>vD</b> If <b>vB</b> value is invalid, the default result is 0b000
Vector Shift Right by Octet	<b>vsro</b>	<b>vD,vA,vB</b>	Shift <b>vA</b> right by the 3 lsbs of <b>vB</b> , and place the result into <b>vD</b> If <b>vB</b> value is invalid, the default result is 0b000

#### 4.2.5.7.1 Immediate Interelement Shifts/Rotates

The Vector Shift Left Double by Octet Immediate (**vsldoi**) instruction provides the basic mechanism that can be used to provide inter-element shifts and/or rotates. This instruction is like a **vperm**, except that the shift count is specified as a literal in the instruction rather than as a control vector in another vector register, as is required by **vperm**. The result vector consists of the left-most 16 bytes of the rotated 32-byte concatenation of **vA:vB**, where shift

(SH) is the rotate count. Table 4-25 below enumerates how various shift functions can be achieved using the **vsidoi** instruction.

**Table 4-25. Coding Various Shifts and Rotates with the vsidoi Instruction**

To Get This:		Code This:			
Operation	sh	Instruction	Immediate	vA	vB
rotate left double	0–15	<b>vsidoi</b>	0–15	MSV	LSV
rotate left double	16–31	<b>vsidoi</b>	mod16(SH)	LSV	MSV
rotate right double	0–15	<b>vsidoi</b>	16–sh	MSV	LSV
rotate right double	16–31	<b>vsidoi</b>	16–mod16(SH)	LSV	MSV
shift left single, zero fill	0–15	<b>vsidoi</b>	0–15	MSV	0x0
shift right single, zero fill	0–15	<b>vsidoi</b>	16–SH	0x0	MSV
rotate left single	0–15	<b>vsidoi</b>	0–15	MSV	=vA
rotate right single	0–15	<b>vsidoi</b>	16–SH	MSV	=vA

#### 4.2.5.7.2 Computed Interelement Shifts/Rotates

The Load Vector for Shift Left (**lvsl**) instruction and Load Vector for Shift Right (**lvsr**) instruction are supplied to assist in shifting and/or rotating vector registers by an amount determined at run time. The input specifications have the same form as the vector load and store instructions, that is, it uses register indirect with index addressing mode(**rA**|0 +**rB**). This is because one of their primary purposes is to compute the permute control vector necessary for post-load and pre-store shifting necessary for dealing with unaligned vectors.

This **lvsl** instruction can be used to align a big-endian unaligned vector after loading the (aligned) vectors that contain its pieces and can be used to unalign a vector register for use in a read-modify-write sequence that will store an unaligned little-endian vector.

The **lvsl** instruction can be used to align a little-endian unaligned vector after loading the (aligned) vectors that contain its pieces and can be used to unalign a vector register for use in a read-modify-write sequence that will store an unaligned big-endian vector.

For an example on how the **lvsl** instruction is used to align a vector in big-endian mode see Section 3.1.6.1, “Accessing a Misaligned Quad Word in Big-Endian Mode.” For an example on how **lvsr** is used to align a vector in little-endian mode see Section 3.1.6.2, “Accessing a Misaligned Quad Word in Little-Endian Mode.”

#### 4.2.5.7.3 Variable Interelement Shifts

A vector register may be shifted left or right by a number of bits specified in a vector register. This operation is supported with four instructions, two for right shift and two for left shift.

The Vector Shift Left by Octet (**vslo**) and Vector Shift Right by Octet (**vsro**) instructions shift a vector register from 0 to 15 bytes as specified in bits 121–124 of another vector register. The Vector Shift Left (**vsl**) and Vector Shift Right (**vsr**) instructions shift a vector register from 0 to 7 bits as specified in another vector register (the shift count must be specified in the three lsbs of each byte in the vector and must be identical in all bytes or the result is boundedly undefined). In all of these instructions, zeros are shifted into vacated element and bit positions.

Used sequentially with the same shift-count vector register, these instructions will shift a vector register left or right from 0 to 127 bits as specified in bits 121–127 of the shift-count vector register. For example:

```
vslo      VZ, VX, VY
vspltb    VY, VY, 15
vsl       VZ, VZ, VY
```

will shift **vX** by the number of bits specified in **vY** and place the results in **vZ**.

With these instructions a full double-register shift can be performed in seven instructions. The following code will shift **vW||vX** left by the number of bits specified in **vY** placing the result in **vZ**:

```
vslo      t1, VW, VY      ; shift the most significant. register left
vspltb    VY, VY, 15
vsl       t1, t1, VY
vsububm   VY, V0, VY      ; adjust count for right shift (V0=0)
vsro      t2, VX, VY      ; right shift least sign. register
vsr       t2, t2, VY
vor       VZ, t1, t2      ; merge to get the final result
```

## 4.2.6 Processor Control Instructions—UISA

Processor control instructions are used to read from and write to the PowerPC condition register (CR), machine state register (MSR), and special-purpose registers (SPRs). See Chapter 4, “Addressing Mode and Instruction Set Summary,” in *PowerPC: The Programming Environments Manual*, for information about the instructions used for reading from and writing to the MSR and SPRs.

### 4.2.6.1 AltiVec Status and Control Register Instructions

Table 4-26 summarizes the instructions for reading from or writing to the AltiVec status and control register (VSCR). For more information on VSCR see section in Section 2.1.1, “The Vector Status and Control Register (VSCR).”

**Table 4-26. Move to/from Condition Register Instructions**

Name	Mnemonic	Syntax	Operation
Move to AltiVec Status and Control Register	<b>mtvscr</b>	CRM,rS	Place the contents of <b>vB</b> into VSCR.

**Table 4-26. Move to/from Condition Register Instructions (Continued)**

Name	Mnemonic	Syntax	Operation
Move from AltiVec Status and Control Register	<b>mfvsr</b>	vB	Place the contents of VSCR into vB.

### 4.2.7 Recommended Simplified Mnemonics

To simplify assembly language programs, a set of simplified mnemonics is provided for some of the most frequently used operations (such as no-op, load immediate, load address, move register, and complement register). Assemblers should provide the simplified mnemonics listed below. Programs written to be portable across the various assemblers for the PowerPC architecture should not assume the existence of mnemonics not described in this document.

Simplified mnemonics are provided for the Data Stream Touch (**dst**) and Data Stream Touch for Store (**dstst**) instructions so that they can be coded with the transient indicator as part of the mnemonic rather than as a numeric operand. Similarly, simplified mnemonics are provided for the Data Stream Stop (**dss**) instruction so that it can be coded with the all streams indicator as part of the mnemonic. These are shown as examples with the instructions.

**Table 4-27. Simplified Mnemonics for Data Stream Touch (dst)**

Operation	Simplified Mnemonic	Equivalent to
Data Stream Touch (non-transient)	<b>dst</b> rA, rB, STRM	<b>dst</b> rA, rB, STRM,0
Data Stream Touch Transient	<b>dstt</b> rA, rB, STRM	<b>dst</b> rA, rB, STRM,1
Data Stream Touch for Store (non-transient)	<b>dstst</b> rA, rB, STRM	<b>dstst</b> rA, rB, STRM,0
Data Stream Touch for Transient	<b>dststt</b> rA, rB, STRM	<b>dststt</b> rA, rB, STRM,1
Data Stream Stop (one stream)	<b>dss</b> STRM	<b>dss</b> STRM,0
Data Stream Stop All	<b>dssall</b>	<b>dss</b> 0,1



## 4.3 AltiVec VEA Instructions

- U The PowerPC virtual environment architecture (VEA) describes the semantics of the memory model that can be assumed by software processes, and includes descriptions of the cache model, cache-control instructions, address aliasing, and other related issues.
- V Implementations that conform to the VEA also adhere to the UISA, but may not necessarily adhere to the OEA. For further details see Chapter 4, “Addressing Mode and Instruction Set Summary,” in *PowerPC: The Programming Environments Manual*.

This section describes the additional instructions that are provided by the AltiVec ISA for the VEA.

### 4.3.1 Memory Control Instructions—VEA

- V Memory control instructions include the following types:
  - - Cache management instructions (user-level and supervisor-level)
    - Segment register manipulation instructions
    - Segment lookaside buffer management instructions
    - Translation lookaside buffer (TLB) management instructions

This section describes the user-level cache management instructions defined by the VEA. See Chapter 4, “Addressing Mode and Instruction Set Summary,” in *PowerPC: The Programming Environments Manual* for more information about supervisor-level cache, segment register manipulation, and TLB management instructions.

#### 4.3.1.1 User-Level Cache Instructions—VEA

- V The instructions summarized in this section provide user-level programs the ability to manage on-chip caches if they are implemented. See Chapter 5, “Cache Model and Memory Coherency,” in *PowerPC: The Programming Environments Manual* for more information about cache topics.

Bandwidth between the processor and memory is managed explicitly by the programmer through the use of cache management instructions. These instructions provide a way for software to communicate to the cache hardware how it should prefetch and prioritize writeback of data. The principle instruction for this purpose is a software directed cache prefetch instruction called Data Stream Touch (**dst**). Other related instructions are provided for complete control of the software directed cache prefetch mechanism.



Table 4-28 summarizes the directed prefetch cache instructions defined by the VEA. Note that these instructions are accessible to user-level programs.



**Table 4-28. User-Level Cache Instructions**

Name	Mnemonic	Syntax	Operation
Data Stream Touch	<b>dst</b>	<b>rA,rB,STRM,T</b>	<p>This instruction associates the data stream specified by the contents of <b>rA</b> and <b>rB</b> with the stream ID specified by <b>STRM</b>.</p> <p>This instruction is a hint that performance will probably be improved if the cache blocks containing the specified data stream are fetched into the data cache, because the program will probably soon load from the stream, and that prefetching from any data stream that was previously associated with the specified stream ID is no longer needed. The hint is ignored for blocks that are Caching Inhibited.</p> <p>The specified data stream is defined by the following.</p> <p>EA: (<b>rA</b>), where <math>rA \wedge 0</math>  unit size: (<b>rB</b>)[3–7] if (<b>rB</b>)[3–7] <math>\wedge 0</math>; otherwise 32  count: (<b>rB</b>)[8–15] if (<b>rB</b>)[8–15] <math>\wedge 0</math>; otherwise 256  stride: (<b>rB</b>)[16–31] if (<b>rB</b>)[16–31] <math>\wedge 0</math>; otherwise 32768</p> <p>The T bit of the instruction indicates whether the data stream is likely to be stored into fairly frequently in the near future (T=0) or to be transient (T=1). If <b>rA</b>=0, the instruction form is invalid.</p>
Data Stream Touch	<b>dstt</b>	<b>rA,rB,STRM,T</b>	<p>This instruction associates the data stream specified by the contents of registers <b>rA</b> and <b>rB</b> with the stream ID specified by <b>STRM</b>.</p> <p>This instruction is a hint that performance will probably be improved if the cache blocks containing the specified data stream are not fetched into the data cache, because the program will probably not load from the stream. That is, the data stream will be relatively transient in nature. That is, it will have poor locality and is likely to be referenced a very few times or over a very short period of time. The memory subsystem can use this persistent/transient knowledge to manage the data as is most appropriate for the specific design of the cache/memory hierarchy of the processor on which the program is executing. An implementation is free to ignore <b>dstt</b>, in that case it should simply be executed as a <b>dst</b>. However, software should always attempt to use the correct form of <b>dst</b> or <b>dstt</b> regardless of whether the intended processor implements <b>dstt</b> or not. In this way the program will automatically benefit when run on processors that do support <b>dstt</b>.</p> <p>The specified data stream is defined by the following.</p> <p>EA: (<b>rA</b>), where <math>rA \wedge 0</math>  unit size: (<b>rB</b>)[3–7] if (<b>rB</b>)[3–7] <math>\wedge 0</math>; otherwise 32  count: (<b>rB</b>)[8–15] if (<b>rB</b>)[8–15] <math>\wedge 0</math>; otherwise 256  stride: (<b>rB</b>)[16–31] if (<b>rB</b>)[16–31] <math>\wedge 0</math>; otherwise 32768</p> <p>The T bit of the instruction indicates whether the data stream is likely to be accessed into fairly frequently in the near future (T=0) or to be transient (T=1). If <b>rA</b>=0, the instruction form is invalid.</p>

**Table 4-28. User-Level Cache Instructions (Continued)**

Name	Mnemonic	Syntax	Operation
Data Stream Touch for Store (non-transient)	<b>dstst</b>	<b>rA,rB,STRM,T</b>	<p>This instruction associates the data stream specified by the contents of registers <b>rA</b> and <b>rB</b> with the stream ID specified by <b>STRM</b>.</p> <p>This instruction is a hint that performance will probably be improved if the cache blocks containing the specified data stream are fetched into the data cache, because the program will probably soon access into the stream, and that prefetching from any data stream that was previously associated with the specified stream ID is no longer needed. The hint is ignored for blocks that are caching inhibited.</p> <p>The specified data stream is defined by the following.</p> <p>EA: (<b>rA</b>), where <math>rA \wedge 0</math>  unit size: (<b>rB</b>)[3–7] if (<b>rB</b>)[3–7] <math>\wedge 0</math>; otherwise 32  count: (<b>rB</b>)[8–15] if (<b>rB</b>)[8–15] <math>\wedge 0</math>; otherwise 256  stride: (<b>rB</b>)[16–31] if (<b>rB</b>)[16–31] <math>\wedge 0</math>; otherwise 32768</p> <p>The T bit of the instruction indicates whether the data stream is likely to be stored into fairly frequently in the near future (T=0) or to be transient (T=1). If <b>rA</b>=0, the instruction form is invalid</p>
Data Stream Touch for Store	<b>dststt</b>	<b>rA,rB,STRM,T</b>	<p>This instruction associates the data stream specified by the contents of <b>rA</b> and <b>rB</b> with the stream ID specified by <b>STRM</b>.</p> <p>This instruction is a hint that performance will probably not be improved if the cache blocks containing the specified data stream are fetched into the data cache, because the program will probably not access the stream. That is, the data stream will be relatively transient in nature. That is, it will have poor locality and is likely to be referenced a very few times or over a very short period of time. The memory subsystem can use this persistent/transient knowledge to manage the data as is most appropriate for the specific design of the cache/memory hierarchy of the processor on which the program is executing.</p> <p>The specified data stream is defined by the following.</p> <p>EA: (<b>rA</b>), where <math>rA \wedge 0</math>  unit size: (<b>rB</b>)[3–7] if (<b>rB</b>)[3–7] <math>\wedge 0</math>; otherwise 32  count: (<b>rB</b>)[8–15] if (<b>rB</b>)[8–15] <math>\wedge 0</math>; otherwise 256  stride: (<b>rB</b>)[16–31] if (<b>rB</b>)[16–31] <math>\wedge 0</math>; otherwise 32768</p> <p>The T bit of the instruction indicates whether the data stream is likely to be stored into fairly frequently in the near future (T=0) or to be transient (T=1). If <b>rA</b>=0, the instruction form is invalid</p>

**Table 4-28. User-Level Cache Instructions (Continued)**

Name	Mnemonic	Syntax	Operation
Data Stream Stop	<b>dss</b>	STRM,A	<p>If A = 0 and a data stream associated with the stream ID specified by STRM exists, this instruction terminates prefetching of that data stream.</p> <p>If A = 1, this instruction terminates prefetching of all existing data streams. (The STRM field is ignored.)</p> <p>In addition, executing a <b>dss</b> instruction ensures that all memory accesses associated with data stream prefetching caused by preceding dst and dstst instructions that specified the same stream ID as that specified by the <b>dss</b> instruction (A = 0), or by all preceding dst and dstst instructions (A = 1), will be in group G1 with respect to the memory barrier created by a subsequent <b>sync</b> instruction.</p> <p><b>dss</b> serves as both a basic and an extended mnemonic. The assembler will recognize a <b>dss</b> mnemonic with two operands as the basic form, and a <b>dss</b> mnemonic with one operand as the extended form.</p> <p>Execution of a <b>dss</b> instruction causes address translation for the specified data stream(s) to cease. Prefetch requests for which the effective address has already been translated may complete and may place the corresponding data into the data cache</p>
Data Stream Stop All	<b>dssall</b>		<p>Terminates prefetching of all existing data streams. All active streams may be stopped.</p> <p>If the optional data stream prefetch facility is implemented, <b>dssall</b> (extended mnemonic for <b>dss</b>), to terminate any data stream prefetching requested by the interrupted program, in order to avoid prefetching data in the wrong context, consuming memory bandwidth fetching data that are not likely to be needed by the other program, and interfering with data cache use by the other program. The <b>dssall</b> must be followed by a <b>sync</b>, and additional software synchronization may be required.</p>

**PRELIMINARY**

# Chapter 5

## Cache, Exceptions, and Memory Management

This chapter summarizes details of the AltiVec™ technology definition that pertain to cache and memory management models. Note that the AltiVec technology defines most of its instructions at the user-level (UISA), so there is little effect on the VEA and OEA portions of the PowerPC architecture definition.

Because the AltiVec Instruction Set Architecture (ISA) uses 128-bit operands, additional instructions are provided to optimize cache and memory bus use.

### 5.1 Memory Bandwidth Management

The AltiVec ISA provides a way for software to speculatively load larger blocks of data from memory. That is, you can use bandwidth that would otherwise be idle which permits the software to take advantage of locality and reduces the number of system memory accesses.

#### 5.1.1 Software-Directed Prefetch

Bandwidth between the processor and memory is managed explicitly by the programmer through use of cache management instructions. These instructions let software indicate to the cache hardware how to prefetch and prioritize writeback of data. The principle instruction for this purpose is a software-directed cache prefetch instruction, Data Stream Touch (**dst**), described in the following section.

##### 5.1.1.1 Data Stream Touch (**dst**)

The data stream prefetch facility permits a program to indicate that a sequence of units of memory is likely to be accessed soon by memory access instructions. Such a sequence is called a data stream or, when the context is clear, simply a stream. A data stream is defined by the following:


- EA—The effective address of the first unit in the sequence
- Unit size—The number of quad words in each unit;  $0 < \text{unit size} \leq 32$
- Count—The number of units in the sequence;  $0 < \text{count} \leq 256$

- **Stride**—The number of bytes between the effective address of one unit in the sequence and the effective address of the next unit in the sequence (that is, the effective address of the  $n$ th unit in the sequence is  $EA + (n - 1) \times \text{stride}$ );  $(-32768 \leq \text{stride} < 0 \text{ or } 0 < \text{stride} \leq 32768)$

The units need not be aligned on a particular memory boundary. The stride may be negative.

The **dst** instruction specifies a starting address, a block size (1–32 vectors), a number of blocks to prefetch (1–256 blocks), and a signed stride in bytes (–32,768 to +32,768 bytes). The 2-bit tag, specified as an immediate field in the opcode, identifies one of four possible touch streams. The starting address of the stream is specified in **rA** (if **rA** = 0, the instruction form is invalid). BlockSize, BlockCount, and BlockStride are specified in **rB**. Do not confuse the term ‘cache block’, the term ‘block’ always indicates a PowerPC cache block.

The format of the **rB** register is shown in Figure 5-1.



**Figure 5-1. Format of rB in dst Instruction**

There is no zero-length block size, block count, or block stride. A BlockSize of 0 indicates 32 vectors, a BlockCount of 0 indicates 256 blocks, and a BlockStride of 0 indicates +32,768 bytes. Otherwise, these fields correspond to the numerical value of the size, count, and stride. Do not specify strides smaller than 1 block (16 bytes).

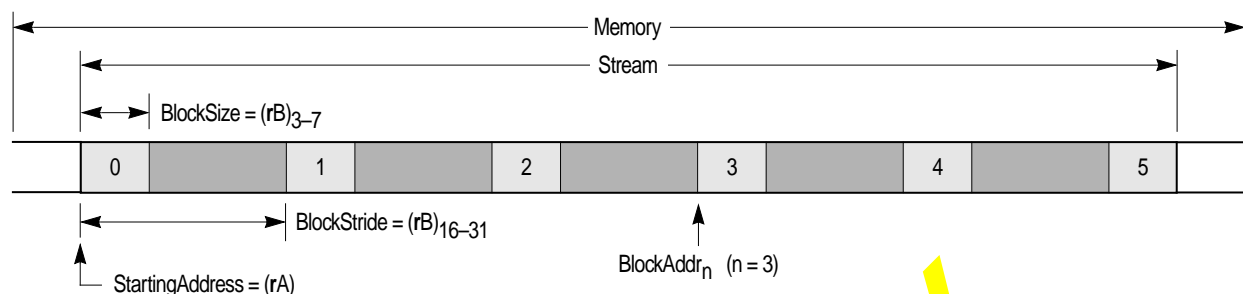
The programmer specifies block size in terms of vectors (16 bytes), regardless of the cache-block size. Hardware automatically optimizes the number of cache blocks it fetches to bring a block into the cache. The number of cache blocks fetched into the cache for each block is the fewest natural cache blocks needed to fetch the entire block, including the effects of block misalignment to cache blocks, as shown in the following:

$$\text{CacheBlocksFetched} = \text{ceiling} \left( \frac{\text{BlockSize} + \text{mod}(\text{BlockAddr}, \text{CacheBlockSize})}{\text{CacheBlockSize}} \right)$$

The address of each block in a stream is a function of the stream’s starting address, the block stride, and the block being fetched. The starting address may be any 32-bit byte address. Each block’s address is computed as a full 32-bit byte address from the following:

$$\text{BlockAddr}_n = (\text{rA}) + n (\text{rB})_{16-31} \quad \begin{array}{l} \text{where } n = \{ 0 \dots (\text{BlockCount} - 1) \} \\ \text{and if } ((\text{rB})_{16-31} = 0) \text{ then } ((\text{rB})_{16-31}) \Leftarrow 32768 \end{array}$$

The address of the first cache block fetched in each block is that block's address aligned to the next lower natural cache-block boundary by ignoring  $\log_2(\text{CacheBlockSize})$  least significant bits (lsbs) (for example, for 32-byte cache-blocks, the five lsbs are ignored). Cache blocks are then fetched sequentially forward until the entire block of vectors is brought into the cache. An example of a six-block data stream is shown in Figure 5-2



**Figure 5-2. Data Stream Touch**

Executing a **dst** instruction notifies the cache/memory subsystem that the program will soon need specified data. If bandwidth is available, the hardware starts loading the specified stream into the cache. To the extent that hardware can acquire the data, when the loads requiring the data finally execute, the target data will be in the cache. Executing a second **dst** to the tag of a stream in progress aborts the existing stream (at hardware's earliest convenience) and establishes a new stream with the same stream tag ID.

The **dst** instruction is a hint to hardware and has no architecturally visible effects (in the PowerPC UISA sense). The hardware is free to ignore it, to start the prefetch when it can, to abort the stream at any time, or to prioritize other memory operations ahead of it. If a stream is aborted, the program still functions properly, but subsequent loads experience the full latency of a cache miss.

The **dst** instruction does not introduce implementation problems like those of load/store multiple/string instructions. Because **dst** does not affect the architectural state, it does not cause interlock problems associated with load/store multiple/string instructions. Also, **dst** does take exceptions and requires no complex recovery mechanism.

Touch instructions should be considered strong hints. Using them in highly speculative situations could waste considerable bandwidth. Implementations that do not implement the stream mechanism treat stream instructions (**dst**, **dstt**, **dsts**, **dstst**, **dss**, and **dssall**) as no-ops. If the stream mechanism is implemented, all four streams must be provided.

### 5.1.1.2 Transient Streams (**dstt**)

The memory subsystem considers **dst** an indication that its stream data is likely to have some reasonable degree of locality and be referenced several times or over some reasonably long period. This is called persistence. The Data Stream Touch Transient instruction (**dstt**) indicates to the memory system that its stream data is transient, that is, it will have poor

locality and is likely to be used very few times or only for a very short time. A memory subsystem can use this knowledge to manage data for the processor's cache/memory design. An implementation may ignore the distinction between transience and persistence, in that case **dstt** acts like **dst**. However, portable software should always use the correct form of **dst** or **dstt** regardless of whether the intended processor makes that distinction.

### 5.1.1.3 Storing to Streams (dstst)

A **dst** instruction brings a cache block into the cache subsystem in a state most efficient for subsequent reading of data from it (load). The companion instruction, Data Stream Touch for Store (**dstst**), brings the cache block into the cache subsystem in a state most efficient for subsequent writing to it (store). For example, in a MESI cache subsystem, a **dst** might bring a cache block in shared (S) state, whereas a **dstst** would bring the cache block in exclusive (E) state to avoid a subsequent demand-driven bus transaction to take ownership of the cache block so the store can proceed.

The **dstst** streams are the same physical streams as **dst** streams, that is, **dstst** stream tags are aliases of **dst** tags. If not implemented, **dstst** defaults to **dst**. If **dst** is not implemented, it is a no-op. The **dststt** instruction is a transient version of **dstst**.

Data stream prefetching of memory locations for which bit 57 of the segment table entry or bit 0 of the segment register (SR) is set is not supported. If a **dst** or **dstst** instruction specifies a data stream containing such memory locations, results are undefined.

### 5.1.1.4 Stopping Streams (dss)

The **dst** instructions have a counterpart called Data Stream Stop (**dss**). A program can stop any given stream prefetch by executing **dss** with that stream's tag. This is useful when a program speculatively starts a stream prefetch but later determines that the instruction stream went the wrong way. The **dss** instruction can stop the stream so no more bandwidth is wasted. All active streams may be stopped by using **dssall**. This is useful when the operating system needs to stop all active streams (process switch) but does not know how many streams are in progress.

Because **dssall** does not specify the number of implemented streams, it should always be used instead of a sequence of **dss** instructions to stop all streams.

Neither **dss** nor **dssall** is execution synchronizing; the time between when a **dss** is issued and the stream stops is not specified. Therefore, when software must ensure that the stream is physically stopped before continuing (for example, before changing virtual memory mapping), a special sequence of synchronizing instructions is required. The sequence can differ for different situations, but the following sequence works in all contexts:

```
dssall ; stop all streams
sync  ; insert a barrier in memory pipe
lwz   Rn,...; stick one more operation in memory pipe
cmpd  Rn,Rn;
bne-  *-4; make sure load data is back
isync ; wait for all previous instructions to complete to ensure
      ; memory pipe is clear and nothing is pending in the old context
```



Data stream prefetching for a given stream is terminated by executing the appropriate **dss** instruction. The termination can be synchronized by executing a **sync** instruction after the **dss** instruction if the memory barrier created by **sync** orders all address translation effects of the subsequent context-altering instructions. Otherwise, data dependencies are also required. For example, the following instruction sequence terminates all data stream prefetching before altering the contents of an segment register (SR):

```
dssall # stop all data stream prefetching
sync  # order dssall before load
lwz   Ry,sr_y(Rx)# load new SR value
mtsr  y,Ry # alter SR y
```

The **mtsr** instruction cannot be executed until the **lwz** loads the SR value into **ry**. The memory access caused by the **lwz** cannot be performed until the **dssall** instruction takes effect (that is, until address translation stops for all data streams and all memory accesses associated with data stream prefetches for which the effective address was translated before the translation stopped are performed).

#### 5.1.1.5 Exception Behavior of Prefetch Streams

In general, exceptions do not cancel streams. Streams are sensitive to whether the processor is in user or supervisor mode (determined by **MSR[PR]**) and whether data address translation is used (determined by **MSR[DR]**). This allows prefetch streams to behave predictably when an exception occurs.

Streams are suspended in real addressing mode (**MSR[DR] = 0**) and remain suspended until translation is turned back on (**MSR[DR]** is set). A **dst** instruction issued while data translation is off (**MSR[DR] = 0**) produces boundedly-undefined results.

A stream is suspended whenever the **MSR[PR]** is different than it was when the **dst** that established it was issued. For example, if a **dst** is issued in user mode (**MSR[PR] = 1**), the resulting stream is suspended when the processor enters supervisor mode (**MSR[PR] = 0**) and remains suspended until the processor returns to user mode. Conversely, if the **dst** were issued in supervisor mode, it is suspended if the machine enters user mode.

Because exceptions do not cancel streams automatically, the operating system must stop streams explicitly when warranted, for example when switching processes or changing virtual memory context.

Take care if data stream prefetching is used in supervisor-level state (**MSR[PR] = 0**).

After an exception, the supervisor-level program that next changes **MSR[DR]** from 0 to 1 cause data-stream prefetching to resume for any data streams for which the corresponding **dst** or **dstst** instruction was executed in supervisor mode; such streams are called supervisor-level data streams. This program is unlikely to be the one that executed the corresponding **dst** or **dstst** instruction and is unlikely to use the same address translation context as that in which the **dst** or **dstst** was executed. (Suspension and resumption of data stream prefetching work more naturally for non-supervisor level data streams, because the next application program to be dispatched after an exception occurs is likely to be the most

recently interrupted program.) Thus, an exception handler that changes the context in which data addresses are translated may need to terminate data-stream prefetching for supervisor-level data streams and to synchronize the termination before changing MSR[DR] to 1.

Although terminating all data stream prefetching in this case would satisfy the requirements of the architecture, doing so would adversely effect the performance of applications that use data-stream prefetching. Thus, it may be better for the operating system to record stream IDs associated with any supervisor-level data streams and to terminate prefetching for those streams only.

Cache effects of supervisor-level data-stream prefetching can also adversely effect performance of applications that use data stream prefetching, as supervisor-level use of the associated stream ID can take over an applications' data stream.

Data stream instructions cannot cause exceptions directly. Therefore, any event that would cause an exception on a normal load or store, such as a page fault or protection violation, is instead aborted and ignored.

Suspension, termination, or supplantation of data stream prefetching for a given data stream need not cancel prefetch requests for that data stream for which the effective address has been translated and need not cause data returned by such requests to be discarded. However, to improve software's ability to pace data stream prefetching with data consumption, it may be better to limit the number of these pending requests that can exist simultaneously.

#### 5.1.1.6 Synchronization Behavior of Streams

Streams are not affected (stopped or suspended) by execution of any PowerPC synchronization instructions (**sync**, **isync**, or **eieio**). This permits these instructions to be used for synchronizing multiple processors without disturbing background prefetch streams. Prefetch streams have no architecturally observable effects and are not affected by synchronization instructions. Synchronizing the termination of data stream prefetching is needed only by the operating system.

#### 5.1.1.7 Address Translation for Streams

Like **dcbt** and **dcbst** instructions, **dst**, **dstst**, **dstt**, and **dststt** are treated as loads with respect to address translation, memory protection, and reference and change recording.

Unlike **dcbt** and **dcbst** instructions, stream instructions that cause a TLB miss cause a page table search and the page descriptor to be loaded into the TLB. Conceptually, address translation and protection checking is performed on every cache-block access in the stream and proceeds normally across page boundaries and TLB misses, terminating only on page faults or protection violations that cause a DSI exception.

Stream instructions operate like normal PowerPC cache instructions (such as **dcbt**) with respect to guarded memory; they are not subject to normal restrictions against prefetching in guarded space because they are program directed. However, **dst** will not start a prefetch stream to guarded space if it itself is speculative.

If the effective address of a cache block within a data stream cannot be translated, or if loading from the block would violate memory protection, the processor will terminate prefetching of that stream. (Continuing to prefetch subsequent cache blocks within the stream might cause prefetching to get too far ahead of consumption of prefetched data.) If the effective address can be translated, a TLB miss will not cause such termination, even on implementations for which TLBs are reloaded in software.

#### 5.1.1.8 Stream Usage Notes

A given data stream exists if a **dst** or **dstst** instruction has been executed that specifies the stream and prefetching of the stream has neither completed, terminated, or been supplanted. Prefetching of the stream has completed, when all the memory locations within the stream that will ever be prefetched as a result of executing the **dst** or **dstst** instruction have been prefetched (for example, locations for which the effective address cannot be translated will never be prefetched). Prefetching of the stream is terminated by executing the appropriate **dss** instruction; it is supplanted by executing another **dst** or **dstst** instruction that specifies the stream ID associated with the given stream. Because there are four stream IDs, as many as four data streams may exist simultaneously.

The maximum block count of **dst** is small because of its preferred usage. It is not intended for a single **dst** instruction to prefetch an entire data stream. Instead, **dst** instructions should be issued periodically, for example on each loop iteration, for the following reasons:

- Short, frequent **dst** instructions better synchronize the stream with consumption.
- With prefetch closely synchronized just ahead of consumption, another activity is less likely to inadvertently evict prefetched data from the cache before it is needed.
- The prefetch stream is restarted automatically after an exception (that could have caused the stream to be terminated by the operating system) with no additional complex hardware mechanisms needed to restart the prefetch stream.

Issuing new **dst** instructions to stream tag IDs in progress terminates old streams—**dst** instructions cannot be queued.

For example, when multiple **dst** instructions are used to prefetch a large stream, it would be poor strategy to issue a second **dst** whose stream begins at the specified end of the first stream before it was certain that the first stream had completed. This could terminate the first stream prematurely, leaving much of the stream unprefetched.

Paradoxically, it would also be unwise to wait for the first stream to complete before issuing the second **dst**. Detecting completion of the first stream is not possible, so the program would have to introduce a pessimistic waiting period before restarting the stream and then incur the full start-up latency of the second stream.

The correct strategy is to issue the second **dst** well before the anticipated completion of the first stream and begin it at an address overlapping the first stream by an amount sufficient to cover any portion of the first stream that could not yet have been prefetched. Issuing the

second **dst** too early is not a concern because blocks prefetched by the first stream hit in the cache and need not be refetched. Thus, even if issued prematurely and overlapped excessively, the second **dst** rapidly advances to the point of prefetching new blocks. This strategy allows a smooth transition from the first stream to the second without significant breaks in the prefetch stream.

For the greatest performance benefit from data-stream prefetching, use the **dst** and **dstst** (and **dss**) instructions so that the prefetched data is used soon after it is available in the data cache. Pacing data stream prefetching with consumption increases the likelihood that prefetched data is not displaced from the cache before it is used, and reduces the likelihood that prefetched data displaces other data needed by the program.

Specifying each logical data stream as a sequence of shorter data streams helps achieve the desired pacing, even in the presence of exceptions, and address translation failures. The components of a given logical data stream should have the following attributes:

- The same stream ID should be associated with each component.
- The components should partially overlap (that is, the first part of a component should consist of the same memory locations as the last part of the preceding component).
- The memory locations which do not overlap with the next component should be large enough that a substantial portion of the component is prefetched. That is, prefetch enough memory locations for the current component before it is taken over by the prefetching being done for the next component.

#### 5.1.1.9 Stream Implementation Assumptions

Some processors can treat **dst** instructions as no-ops. However, if a processor implements **dst**, a minimum level of functionality will be provided to create as consistent a programming model across different machines as possible. Programs can assume the following:

- Implements all four tagged streams
- Implements each tagged stream as a separate, independent stream with arbitration for memory access performed on a round-robin basis.
- Searches the table for each stream access that misses in the TLB.
- Does not abort streams on page boundary crossings
- Does not abort streams on exceptions (except DSI exceptions caused by the stream).
- Does not abort streams, or hold up execution pending completion of streams, on the PowerPC synchronization instructions **sync**, **isync**, or **eieio**.
- Does not abort streams on TLB misses that occur on loads or stores issued concurrently with running streams. However, a DSI exception from one of those loads or stores may cause streams to abort.

### 5.1.2 Prioritizing Cache Block Replacement

Load Vector Indexed LRU (**lvxl**) and Store Vector Indexed LRU (**stvx**) instructions provide explicit control over cache block replacement by letting the programmer indicate whether an access is likely to be the last reference made to the cache block containing this load or store. The cache hardware can then prioritize replacement of this cache block over others with older but more useful data.

Data accessed by a normal load or store is likely to be needed more than once. Marking this data as most-recently used (MRU) indicates that it should be a low-priority candidate for replacement. However, some data, such as that used in DSP multimedia algorithms, is rarely reused and should be marked as the highest priority candidate for replacement.

Normal accesses mark data MRU. Data unlikely to be reused can be marked LRU. For example, on replacing a cache block marked LRU by one of these instructions, a processor may improve cache performance by evicting the cache block without storing it in intermediate levels of the cache hierarchy (except to maintain cache consistency).

#### 5.1.2.1 Partially Executed AltiVec Instructions

The OEA permits certain instructions to be partially executed when an alignment or DSI exception occurs. In the same way that the target register may be altered when floating-point load instructions cause a DSI exception, if the AltiVec facility is implemented, the target register (**vD**) may be altered when **lvx** or **lvxl** is executed and the TLB entry is invalidated part way through the access.

Exceptions cause data stream prefetching to be suspended for all existing data streams. Prefetching for a given data stream resumes when control is returned to the interrupted program, if the stream still exists (for example, the operating system did not terminate prefetching for the stream).

## 5.2 DSI Exception—Data Address Breakpoint

A DABR match causes a DSI exception in implementations that support the data breakpoint feature. When an DABR match occurs on a non-AltiVec PowerPC processor, the DAR is set to any effective address between and including the word (for a byte, half word, or word access) or double word (for a double-word access) specified by the effective address computed by the instruction and the effective address of the last byte in the word or double word in which the match occurred. In processors that support AltiVec technology, this would include a quad word access from an **lvx**, **lvxl**, **stvx**, or **stvx** instruction to a segment or BAT area.

## 5.3 AltiVec Unavailable Exception (0x00F20)

The AltiVec facility includes an additional instruction-caused, precise exception to those defined by the OEA and discussed in Chapter 6, “Exceptions,” in the *PowerPC Programming Environments Manual*. An AltiVec unavailable exception occurs when no

higher priority exception exists (see Table 5-2), an attempt is made to execute an AltiVec instruction, and MSR[VEC] = 0.

Register settings for AltiVec unavailable exceptions are described in Table 5-1.

**Table 5-1. AltiVec Unavailable Exception—Register Settings**

Register	Setting Description																				
SRR0	Set to the effective address of the instruction that caused the exception																				
SRR1	<div><div>32-Bit</div><div><div>0</div><div>Loaded with equivalent bits from the MSR</div></div><div><div>1–4</div><div>Cleared</div></div><div><div>5–9</div><div>Loaded with equivalent bits from the MSR</div></div><div><div>10–15</div><div>Cleared</div></div><div><div>16–31</div><div>Loaded with equivalent bits from the MSR</div></div><div>Note that depending on the implementation, additional bits in the MSR may be copied to SRR1.</div></div>																				
MSR	<table><tr><td>SF<sup>1</sup> 1</td><td>EE 0</td><td>SE 0</td><td>DR 0</td></tr><tr><td>ISF<sup>1</sup> —</td><td>PR 0</td><td>BE 0</td><td>RI 0</td></tr><tr><td>VEC 0</td><td>FP 0</td><td>FE1 0</td><td>LE Set to value of ILE</td></tr><tr><td>POW 0</td><td>ME —</td><td>IP —</td><td></td></tr><tr><td>ILE —</td><td>FE0 0</td><td>IR 0</td><td></td></tr></table>	SF <sup>1</sup> 1	EE 0	SE 0	DR 0	ISF <sup>1</sup> —	PR 0	BE 0	RI 0	VEC 0	FP 0	FE1 0	LE Set to value of ILE	POW 0	ME —	IP —		ILE —	FE0 0	IR 0	
SF <sup>1</sup> 1	EE 0	SE 0	DR 0																		
ISF <sup>1</sup> —	PR 0	BE 0	RI 0																		
VEC 0	FP 0	FE1 0	LE Set to value of ILE																		
POW 0	ME —	IP —																			
ILE —	FE0 0	IR 0																			

<sup>1</sup> 64-bit implementations only

When an AltiVec unavailable exception is taken, instruction execution resumes as offset 0x00F20 from the base address determined by MSR[IP].

The **dst** and **dstst** instructions are supported if MSR[DR] = 1. If either instruction is executed when MSR[DR] = 0 (real addressing mode), results are boundedly undefined.

Conditions that cause this exception are prioritized among instruction-caused (synchronous), precise exceptions as shown in Table 5-2 (taken from the section “Exception Priorities,” in Chapter 6, “Exceptions,” in *PowerPC: The Programming Environments Manual*).



**Table 5-2. Exception Priorities (Synchronous/Precise Exceptions)**

Priority	Exception
3 <sup>1</sup>	<p>Instruction dependent—When an instruction causes an exception, the exception mechanism waits for any instructions prior to the excepting instruction in the instruction stream to complete. Any exceptions caused by these instructions are handled first. It then generates the appropriate exception if no higher priority exception exists when the exception is to be generated.</p> <p>Note that a single instruction can cause multiple exceptions. When this occurs, those exceptions are ordered in priority as indicated in the following:</p> <ul style="list-style-type: none"> <li>A. Integer loads and stores <ul style="list-style-type: none"> <li>a. Alignment</li> <li>b. DSI</li> <li>c. Trace (if implemented)</li> </ul> </li> <li>B. Floating-point loads and stores <ul style="list-style-type: none"> <li>a. Floating-point unavailable</li> <li>b. Alignment</li> <li>c. DSI</li> <li>d. Trace (if implemented)</li> </ul> </li> <li>C. Other floating-point instructions <ul style="list-style-type: none"> <li>a. Floating-point unavailable</li> <li>b. Program—Precise-mode floating-point enabled exception</li> <li>c. Floating-point assist (if implemented)</li> <li>d. Trace (if implemented)</li> </ul> </li> <li>D. AltiVec Loads and Stores (if AltiVec facility implemented) <ul style="list-style-type: none"> <li>a. AltiVec unavailable</li> <li>b. DSI</li> <li>c. Trace (if implemented)</li> </ul> </li> <li>E. Other AltiVec Instructions (if AltiVec facility implemented) <ul style="list-style-type: none"> <li>a. AltiVec unavailable</li> <li>b. Trace (if implemented)</li> </ul> </li> <li>F. <b>rfi</b> and <b>mtmsr</b> <ul style="list-style-type: none"> <li>a. Program—Supervisor level Instruction</li> <li>b. Program—Precise-mode floating-point enabled exception</li> <li>c. Trace (if implemented), for <b>mtmsr</b> only</li> </ul> <p>If precise-mode IEEE floating-point enabled exceptions are enabled and the FPSCR[FEX] bit is set, a program exception occurs no later than the next synchronizing event.</p> </li> <li>G. Other instructions <ul style="list-style-type: none"> <li>a. These exceptions are mutually exclusive and have the same priority: <ul style="list-style-type: none"> <li>— Program: Trap</li> <li>— System call (<b>sc</b>)</li> <li>— Program: Supervisor level instruction</li> <li>— Program: Illegal Instruction</li> </ul> </li> <li>b. Trace (if implemented)</li> </ul> </li> <li>F. ISI exception</li> </ul> <p>The ISI exception has the lowest priority in this category. It is only recognized when all instructions prior to the instruction causing this exception appear to have completed and that instruction is to be executed. The priority of this exception is specified for completeness and to ensure that it is not given more favorable treatment. An implementation can treat this exception as though it had a lower priority.</p>

<sup>1</sup> The exceptions are third in priority after system reset and machine check exceptions

**PRELIMINARY**



# Chapter 6

## AltiVec Instructions

This chapter lists the AltiVec instruction set in alphabetical order by mnemonic. Note that each entry includes the instruction format and a graphical representation of the instruction. All the instructions are 32 bit and a description of the instruction fields and pseudocode conventions are also provided. For more information on the AltiVec instruction set, refer to Chapter 4 “Addressing Modes and Instruction Set Summary,” For more information on the PowerPC instruction set, refer to Chapter 8, “Instruction Set,” in *The PowerPC Microprocessor Family: The Programming Environments Manual*.

### 6.1 Instruction Formats

AltiVec instructions are four bytes (32 bits) long and are word-aligned. AltiVec instruction set architecture (ISA) has 4 operands, three source vectors and one result vector. Bits 0–5 always specify the primary opcode for AltiVec instructions. AltiVec ALU type instructions specify the primary opcode point 4 (0b000100). AltiVec load, store, and stream prefetch instructions use secondary opcodes in primary opcode 31 (0b011111).

Within a vector register, a byte, half word or word element, are referred to as follows:

- Byte elements, each byte = 8 bits, so in the pseudocode,  $n = 8$  and there would be a total of 16 elements
- Half-word elements, each byte = 16 bits, so in the pseudocode,  $n = 16$  and there would be a total of 8 elements
- Word elements, each byte = 32bits, so in the pseudocode,  $n = 32$  and there would be a total of 4 elements

Refer to Figure 1-3, for an example of how elements are placed in a vector register.

## 6.1.1 Instruction Fields

Table 6-1 describes the instruction fields used in the various instruction formats.

**Table 6-1. Instruction Syntax Conventions**

Field	Description
OPCD (0–5)	Primary opcode field
rA, A(11–15)	Specifies a GPR to be used as a source or destination.
rB, B(16–20)	Specifies a GPR to be used as a source.
Rc (31)	Record bit. 0 Does not update the condition register (CR). 1 For the optional AltiVec facility, set CR field 6 to control program flow as described in Chapter 2, Section 2.1.3 “Condition register”
vA (11–15)	Specifies a vector register to be used as a source
vB (16–20)	Specifies a vector register to be used as a source.
vC (21–25)	Specifies a vector register to be used as a source.
vD (6–10)	Specifies a vector register to be used as a destination.
vS (6–10)	Specifies a vector register in the VRF to be used as a source.
SHB (22–25)	Specifies a shift amount in bytes.
SIMM (11–15)	This immediate field is used to specify a (5 bit) signed integer.
tag (9-10)	
UIMM (11–15)	This immediate field is used to specify a 4,8,12, or 16-bit unsigned integer.
XO	Extended Opcode Field.

## 6.1.2 Notation and Conventions

The operation of some instructions is described by a semiformal language (pseudocode). See Table 6-2 for a list of additional pseudocode notation and conventions used throughout this section.

**Table 6-2. Notation and Conventions**

Notation/Convention	Meaning
$\leftarrow$	Assignment
$\neg$	NOT logical operator
do i=X,Y,Z	do the following starting at X and iterating to Y by Z
$+_{\text{int}}$	2's complement integer add
$-_{\text{int}}$	2's complement integer subtract
$+_{\text{ui}}$	Unsigned integer add
$-_{\text{ui}}$	Unsigned integer subtract

**Table 6-2. Notation and Conventions (Continued)**

Notation/Convention	Meaning
* <sub>ui</sub>	Unsigned integer multiply
+ <sub>si</sub>	Signed integer add
- <sub>si</sub>	Signed integer subtract
* <sub>si</sub>	Signed integer multiply
* <sub>sui</sub>	Signed integer (first operand) multiplied by unsigned integer (second operand) producing signed result
/	Integer divide
+ <sub>fp</sub>	Single-precision floating-point add
- <sub>fp</sub>	Single-precision floating-point subtract
* <sub>fp</sub>	Single-precision floating-point multiply
÷ <sub>fp</sub>	Single-precision floating-point divide
< <sub>ui</sub> , ≤ <sub>ui</sub> , > <sub>ui</sub> , ≥ <sub>ui</sub>	Unsigned integer comparison relations
< <sub>si</sub> , ≤ <sub>si</sub> , > <sub>si</sub> , ≥ <sub>si</sub>	Signed integer comparison relations
< <sub>fp</sub> , ≤ <sub>fp</sub> , > <sub>fp</sub> , ≥ <sub>fp</sub>	Single precision floating point comparison relations
^=	Not equal
= <sub>int</sub>	Integer equal to
= <sub>ui</sub>	Unsigned integer equal to
= <sub>si</sub>	Signed integer equal to
= <sub>fp</sub>	floating-point equal to
X >> <sub>ui</sub> Y	shift X right by Y bits extending Xs vacated bits with zeros
X >> <sub>si</sub> Y	shift X right by Y bits extending Xs vacated bits with the sign bit of X
X << <sub>ui</sub> Y	shift X left by Y bits inserting Xs vacated bits with zeros
	Used to describe the concatenation of two values (that is, 010    111 is the same as 010111)
&,	AND logical operator
	OR logical operator
?	undefined (see Load Vector)
X <sub>0</sub>	X zeros
X <sub>1</sub>	X ones
X <sub>Y</sub>	X copies of Y
X <sub>Y</sub>	bit Y of X
X <sub>Y:Z</sub>	bits Y through Z, inclusive, of X
LENGTH(x)	Length of x, in bits. If x is the word “element”, LENGTH(x) is the length, in bits, of the element implied by the instruction mnemonic.

**Table 6-2. Notation and Conventions (Continued)**

<b>Notation/Convention</b>	<b>Meaning</b>
ROTL(x,y)	Result of rotating x left by y bits
UtoUImod(X,Y)	Chop unsigned integer X- to Y-bit unsigned integer
UtoUlsat(X,Y)	Result of converting the unsigned-integer x to a y-bit unsigned-integer with unsigned-integer saturation
SltUlsat(X,Y)	Result of converting the signed-integer x to a y-bit unsigned-integer with unsigned-integer saturation
SltSImod(X,Y)	Chop integer X- to Y-bit integer
SltSlsat(X,Y)	Result of converting the signed-integer x to a y-bit signed-integer with signed-integer saturation
RndToNearFP32	The single-precision floating-point number that is nearest in value to the infinitely-precise floating-point intermediate result x (in case of a tie, the even single-precision floating-point value is used)
RndToFPInt32Near	The value x if x is a single-precision floating-point integer; otherwise the single-precision floating-point integer that is nearest in value to x (in case of a tie, the even single-precision floating-point integer is used)
RndToFPInt32Trunc	The value x if x is a single-precision floating-point integer; otherwise the largest single-precision floating-point integer that is less than x if x>0, or the smallest single-precision floating-point integer that is greater than x if x<0
RndToFPInt32Ceil	The value x if x is a single-precision floating-point integer; otherwise the smallest single-precision floating-point integer that is greater than x
RndToFPInt32Floor	The value x if x is a single-precision floating-point integer; otherwise the largest single-precision floating-point integer that is less than x
CnvtFP32ToUI32Sat(x)	Result of converting the single-precision floating-point value x to a 32-bit unsigned-integer with unsigned-integer saturation
CnvtFP32ToSI32Sat(x)	Result of converting the single-precision floating-point value x to a 32-bit signed-integer with signed-integer saturation
CnvtUI32ToFP32(x)	Result of converting the 32-bit unsigned-integer x to floating-point single format
CnvtSI32ToFP32(x)	Result of converting the 32-bit signed-integer x to floating-point single format
MEM(X,Y)	Value at memory location X of size Y bytes
SwapDouble	swap the doublewords in a quadword vector
ZeroExtend(X,Y)	zero-extend X on the left with zeros to produce Y-bit value
SignExtend(X,Y)	sign-extend X on the left with sign bits (that is, with copies of bit 0 of x) to produce Y-bit value
RotateLeft(X,Y)	rotate X left by Y bits
mod(X,Y)	remainder of X/Y
UImaximum(X,Y)	maximum of 2 unsigned integer values, X and Y
SImaximum(X,Y)	maximum of 2 signed integer values, X and Y
FPmaximum(X,Y)	maximum of 2 floating-point values, X and Y

**Table 6-2. Notation and Conventions (Continued)**

Notation/Convention	Meaning
UIminimum(X,Y)	minimum of 2 unsigned integer values, X and Y
SIminimum(X,Y)	minimum of 2 unsigned integer values, X and Y
FPminimum(X,Y)	minimum of 2 floating-point values, X and Y
FPReciprocalEstimate12(X)	12-bit-accurate floating-point estimate of 1/X
FPReciprocalSQRTEstimate12(X)	12-bit-accurate floating-point estimate of 1/(sqrt(X))
FPLog <sub>2</sub> Estimate3(X)	3-bit-accurate floating-point estimate of log <sub>2</sub> (X)
FPPower2Estimate3(X)	3-bit-accurate floating-point estimate of 2**X
CarryOut(X + Y)	carry out of the sum of X and Y
ROTL[64](x, y)	Result of rotating the 64-bit value x left y positions
ROTL[32](x, y)	Result of rotating the 32-bit value x    x left y positions, where x is 32 bits long
0bnnnn	A number expressed in binary format.
0xnnnn	A number expressed in hexadecimal format.
(n)x	The replication of x, n times (that is, x concatenated to itself n – 1 times). (n)0 and (n)1 are special cases. A description of the special cases follows: <ul style="list-style-type: none"> <li>• (n)0 means a field of n bits with each bit equal to 0. Thus (5)0 is equivalent to 0b00000.</li> <li>• (n)1 means a field of n bits with each bit equal to 1. Thus (5)1 is equivalent to 0b11111.</li> </ul>
(rA 0)	The contents of rA if the rA field has the value 1–31, or the value 0 if the rA field is 0.
(rX)	The contents of rX
x[n]	n is a bit or field within x, where x is a register
x <sup>n</sup>	x is raised to the nth power
ABS(x)	Absolute value of x
CEIL(x)	Least integer ≥ x
Characterization	Reference to the setting of status bits in a standard way that is explained in the text.
CIA	Current instruction address. The 32-bit address of the instruction being described by a sequence of pseudocode. Used by relative branches to set the next instruction address (NIA) and by branch instructions with LK = 1 to set the link register. Does not correspond to any architected register.
Clear	Clear the leftmost or rightmost n bits of a register to 0. This operation is used for rotate and shift instructions.
Clear left and shift left	Clear the leftmost b bits of a register, then shift the register left by n bits. This operation can be used to scale a known non-negative array index by the width of an element. These operations are used for rotate and shift instructions.
Cleared	Bits = 0.

**Table 6-2. Notation and Conventions (Continued)**

Notation/Convention	Meaning
Do	Do loop. <ul style="list-style-type: none"> <li>• Indenting shows range.</li> <li>• “To” and/or “by” clauses specify incrementing an iteration variable.</li> <li>• “While” clauses give termination conditions.</li> </ul>
DOUBLE(x)	Result of converting x from floating-point single-precision format to floating-point double-precision format.
Extract	Select a field of <i>n</i> bits starting at bit position <i>b</i> in the source register, right or left justify this field in the target register, and clear all other bits of the target register to zero. This operation is used for rotate and shift instructions.
EXTS(x)	Result of extending x on the left with sign bits
GPR(x)	General-purpose register x
if...then...else...	Conditional execution, indenting shows range, else is optional.
Insert	Select a field of <i>n</i> bits in the source register, insert this field starting at bit position <i>b</i> of the target register, and leave other bits of the target register unchanged. (No simplified mnemonic is provided for insertion of a field when operating on double words; such an insertion requires more than one instruction.) This operation is used for rotate and shift instructions. (Note that simplified mnemonics are referred to as extended mnemonics in the architecture specification.)
Leave	Leave innermost do loop, or the do loop described in leave statement.
MASK(x, y)	Mask having ones in positions x through y (wrapping if x > y) and zeros elsewhere.
MEM(x, y)	Contents of y bytes of memory starting at address x.
NIA	Next instruction address, which is the 32-bit address of the next instruction to be executed (the branch destination) after a successful branch. In pseudocode, a successful branch is indicated by assigning a value to NIA. For instructions which do not branch, the next instruction address is CIA + 4. Does not correspond to any architected register.
OEA	PowerPC operating environment architecture
Rotate	Rotate the contents of a register right or left <i>n</i> bits without masking. This operation is used for rotate and shift instructions.
ROTL[64](x, y)	Result of rotating the 64-bit value x left y positions
ROTL[32](x, y)	Result of rotating the 64-bit value x    x left y positions, where x is 32 bits long
Set	Bits are set to 1.
Shift	Shift the contents of a register right or left <i>n</i> bits, clearing vacated bits (logical shift). This operation is used for rotate and shift instructions.
SINGLE(x)	Result of converting x from floating-point double-precision format to floating-point single-precision format.
SPR(x)	Special-purpose register x
TRAP	Invoke the system trap handler.
Undefined	An undefined value. The value may vary from one implementation to another, and from one execution to another on the same implementation.

**Table 6-2. Notation and Conventions (Continued)**

<b>Notation/Convention</b>	<b>Meaning</b>
UISA	PowerPC user instruction set architecture
VEA	PowerPC virtual environment architecture

**PRELIMINARY**

Table 6-3 describes instruction field notation conventions used throughout this chapter.

**Table 6-3. Instruction Field Conventions**

The Architecture Specification	Equivalent to:
D	d
DS	ds
FLM	FM
RA, RB, RT, RS	rA, rB, rD, rS
RA, RB, RT, RS	A, B, D, S
SI	SIMM
U	IMM
UI	UIMM
VA, VB, VC, VT, VS	vA, vB, vC, vD, vS
/, //, ///	0...0 (shaded)

Precedence rules for pseudocode operators are summarized in Table 6-4.

**Table 6-4. Precedence Rules**

Operators	Associativity
$x[n]$ , function evaluation	Left to right
$(n)x$ or replication, $x(n)$ or exponentiation	Right to left
unary $-$ , $\neg$	Right to left
$*$ , $\div$	Left to right
$+$ , $-$	Left to right
$\parallel$	Left to right
$=$ , $\neq$ , $<$ , $\leq$ , $>$ , $\geq$ , $<U$ , $>U$ , $?$	Left to right
$\&$ , $\oplus$ , $\equiv$	Left to right
$ $	Left to right
$-$ (range)	None
$\leftarrow$ , $\leftarrow_{ica}$	None



Operators higher in Table 6-4 are applied before those lower in the table. Operators at the same level in the table associate from left to right, from right to left, or not at all, as shown. For example, ‘-’ (unary minus) associates from left to right, so  $a - b - c = (a - b) - c$ . Parentheses are used to override the evaluation order implied by Table 6-4, or to increase clarity; parenthesized expressions are evaluated before serving as operands.


## 6.2 AltiVec Instruction Set

The remainder of this chapter lists and describes the instruction set for the AltiVec architecture. The instructions are listed in alphabetical order by mnemonic. Table 6-5 shows the format for each instruction description page.

Instruction name	<b>vaddsbs</b> Vector Add Signed Byte Saturate	<b>vaddsbs</b>
Instruction syntax	<b>vaddsbs</b> <b>vD,vA,vB</b>	
Instruction encoding		
Pseudocode description of instruction operation	<pre> n ← LENGTH(element) do i=0 to 127 by n   aop0:n ← SignExtend((vA)<sub>i:i+n-1</sub>,n+1)   bop0:n ← SignExtend((vB)<sub>i:i+n-1</sub>,n+1)   temp0:n ← aop0:n +int bop0:n   vD<sub>i:i+n-1</sub> ← SToSIsat(temp0:n,n) </pre>	
Text description of instruction operation	<p>For <b>vaddsbs</b> each element (16 total elements) is a byte (<math>n=8</math>).</p> <p>Each signed-integer element in <b>vA</b> is added to the corresponding signed-integer element <b>vB</b>.</p> <p>If the intermediate result is greater than <math>2^{n-1}-1</math> it saturates to <math>2^{n-1}-1</math> and if it is less <math>-2^{n-1}</math> it saturates to <math>-2^{n-1}</math>, where <math>n</math> is the length of the element.</p> <p>The signed-integer result is placed into the corresponding element of <b>VD</b>.</p>	
Figure showing how elements are placed		
Quick reference legend		

**Table 6-5. Instruction Description**

**dss** STRM,A



$\text{DataStreamPrefetchControl} \leftarrow \text{"stop"} \parallel \text{STRM}$

Note that A does not represent **rA** in this instruction. If A=0 and a data stream associated with the stream ID specified by STRM exists, this instruction terminates prefetching of that data stream. If A=1, this instruction terminates prefetching of all existing data streams (the STRM field is ignored.)


In addition, executing a **dss** instruction ensures that all accesses associated with data stream prefetching caused by preceding **dst** and **dstst** instructions that specified the same stream ID as that specified by the **dss** instruction (A=0), or by all preceding **dst** and **dstst** instructions (A=1), will be in group G1 with respect to the memory barrier created by a subsequent **sync** instruction (see Section 5.1.1.1, “Software-Directed Prefetch,” for more details.)

Other registers altered:

- None

Simplified mnemonics:

**dss** **rA**, **rB**, STRM equivalent to **dss** STRM,0



**Figure 6-1. Data Stream Touch**

dssalldssall

Data Stream Stop All

**dssall** STRM,A

31					A	00		STRM	00000					00000					822					0			
0		5		6	7	8	9	10	11	15					16	20					21	30					31

$$\text{DataStreamPrefetchControl} \leftarrow \text{"stop"} \parallel \text{STRM}$$

Note that A does not represent **rA** in this instruction.

If A=0 and a data stream associated with the stream ID specified by STRM exists, this instruction terminates prefetching of that data stream.

If A=1, this instruction terminates prefetching of all existing data streams. (The STRM field is ignored.)

In addition, executing a **dss** instruction ensures that all accesses associated with data stream prefetching caused by preceding **dst** and **dstst** instructions that specified the same stream ID as that specified by the **dss** instruction (A=0), or by all preceding **dst** and **dstst** instructions (A=1), will be in group G1 with respect to the memory barrier created by a subsequent **sync** instruction (see Section 5.1.1.1, “Software-Directed Prefetch,” for more details).

Other registers altered:

- None

Simplified mnemonics:


**dssall** STRM equivalent to **dss** 0,1

# dst

Data Stream Touch

# dst

**dst**                      **rA,rB,STRM,T**



$\text{addr}_{0:63} \leftarrow (\mathbf{rA})$

$\text{DataStreamPrefetchControl} \leftarrow \text{"start"} \parallel \text{STRM} \parallel \text{static} \parallel (\mathbf{rArB}) \parallel \text{addr}$

This instruction associates the data stream specified by the contents of **rA** and **rB** with the stream ID specified by **STRM**. The specified data stream is defined by the following:

- EA: (**rA**), where  $\mathbf{rA} \wedge = 0$
- Unit size: (**rB**)[35-39 {3-7}] if (**rB**)[35-39 {3-7}]  $\wedge = 0$ ; otherwise 32
- Count: (**rB**)[40-47 {8-15}] if (**rB**)[40-47 {8-15}]  $\wedge = 0$ ; otherwise 256
- Stride: (**rB**)[48-63 {16-31}] if (**rB**)[48-63 {16-31}]  $\wedge = 0$ ; otherwise 32768

The T bit of the instruction indicates whether the data stream is likely to be loaded from fairly frequently in the near future (T = 0) or to be transient (T = 1).

The **dst** instruction does the following:

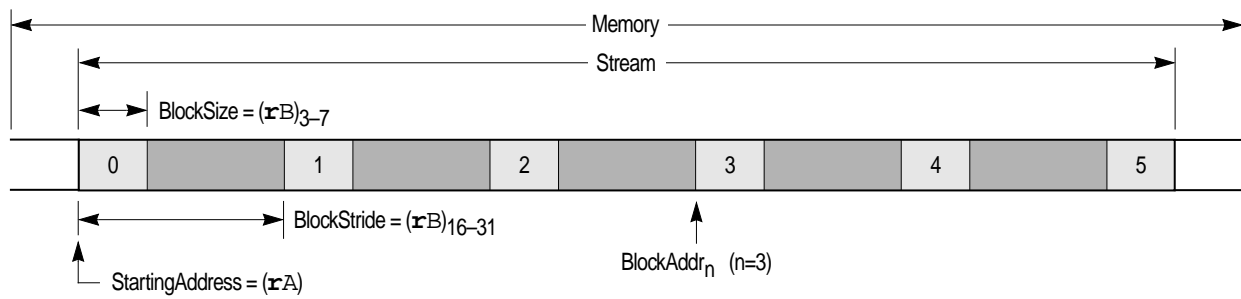
- Indicates that the specified stream may soon be loaded from (**dst**)
- Indicates whether memory locations within the stream are likely to be needed over a longer period of time and should not be treated as transient data
- Associates the stream with the specified stream ID (a number in the range 0–3)
- Indicates that prefetching from any stream that was previously associated with the specified stream ID is no longer needed

Other registers altered:

- None

Simplified mnemonics:

**dst**    **rA, rB, STRM**                      equivalent to                      **dst**    **rA, rB, STRM,0**



**Figure 6-2. Data Stream Touch**


**PRELIMINARY**

# dstst

Data Stream Touch for Store

# dstst

**dstst** **rA,rB,tag**



$\text{addr}_{0-63} \leftarrow (\text{rA})$

$\text{DataStreamPrefetchControl} \leftarrow \text{"start"} \parallel \text{tag} \parallel \text{static} \parallel (\text{rB}) \parallel \text{addr}$

The **dst** instruction does the following:

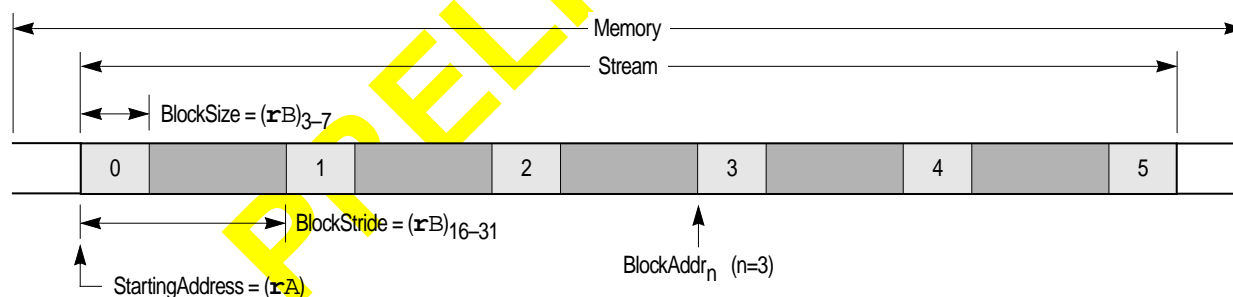
- Indicates that the specified stream may soon be stored to memory (**dstst**)
- Indicates whether memory locations within the stream are likely to be needed over a longer period of time and should not be treated as transient data
- Associates the stream with the specified stream ID (a number in the range 0–3)
- Indicates that prefetching from any stream that was previously associated with the specified stream ID is no longer needed

Other registers altered:

- None

Simplified mnemonics:

**dstst rA, rB, STRM** equivalent to **dstst rA, rB, STRM,0**



**Figure 6-3. Data Stream Touch**

# dststt

Data Stream Touch for Store for Transient

# dststt

**dststt** **rA,rB,tag**

31			1	0 0 0			tag	A			B			11			22			0										
0			5			6	7	9			10	11			15			16	20			21	25			26	30			31

$\text{addr}_{0-63} \leftarrow (\text{rA})$

$\text{DataStreamPrefetchControl} \leftarrow \text{"start"} \parallel \text{tag} \parallel \text{transient} \parallel (\text{rB}) \parallel \text{addr}$

The **dststt** instruction does the following:


- Indicates that the specified stream may soon be accessed to memory (**dststt**).
- Indicates whether memory locations within the stream are likely to be needed during a relatively short period of time and should be treated as transient data.
- Associates the stream with the specified stream ID (a number in the range 0–3)
- Indicates that prefetching from any stream that was previously associated with the specified stream ID is no longer needed

Other registers altered:

- None

Simplified mnemonics:

**dststt rA, rB, STRM** equivalent to **dststt rA, rB, STRM,1**




**Figure 6-4. Data Stream Touch**

# dstt

Data Stream Touch Transient

# dstt

**dstt** **rA,rB,tag**



$\text{addr}_{0-63} \leftarrow (\text{rA})$

$\text{DataStreamPrefetchControl} \leftarrow \text{"start"} \parallel \text{tag} \parallel \text{transient} \parallel (\text{rB}) \parallel \text{addr}$

The **dstt** instruction does the following:

- Indicates that the specified stream may soon be loaded from memory (**dstt**).
- Indicates whether memory locations within the stream are likely to be needed during a relatively short period of time and should be treated as transient data.
- Associates the stream with the specified stream ID (a number in the range 0–3)
- Indicates that prefetching from any stream that was previously associated with the specified stream ID is no longer needed

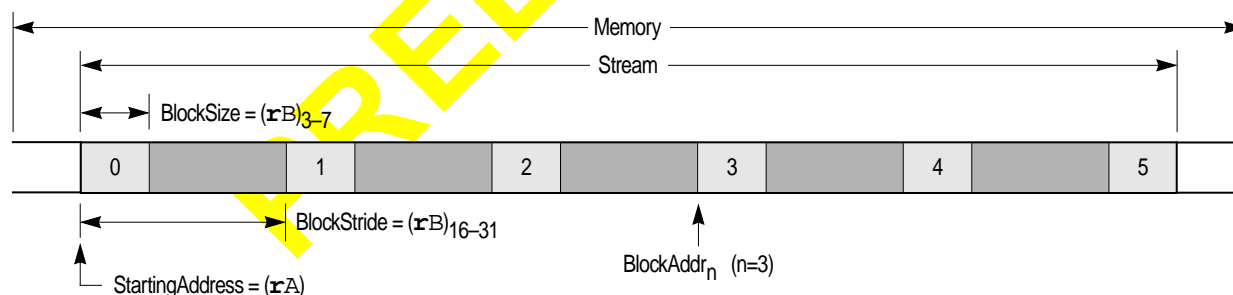
Other registers altered:

- None

Simplified mnemonics:

**dstt** **rA, rB, STRM** equivalent to **dstt** **rA, rB, STRM,1**

**Figure 6-5**




**Figure 6-6. Data Stream Touch**



**PRELIMINARY**

Load Vector Element Byte Indexed

**lvebx** **vD,rA,rB**

```

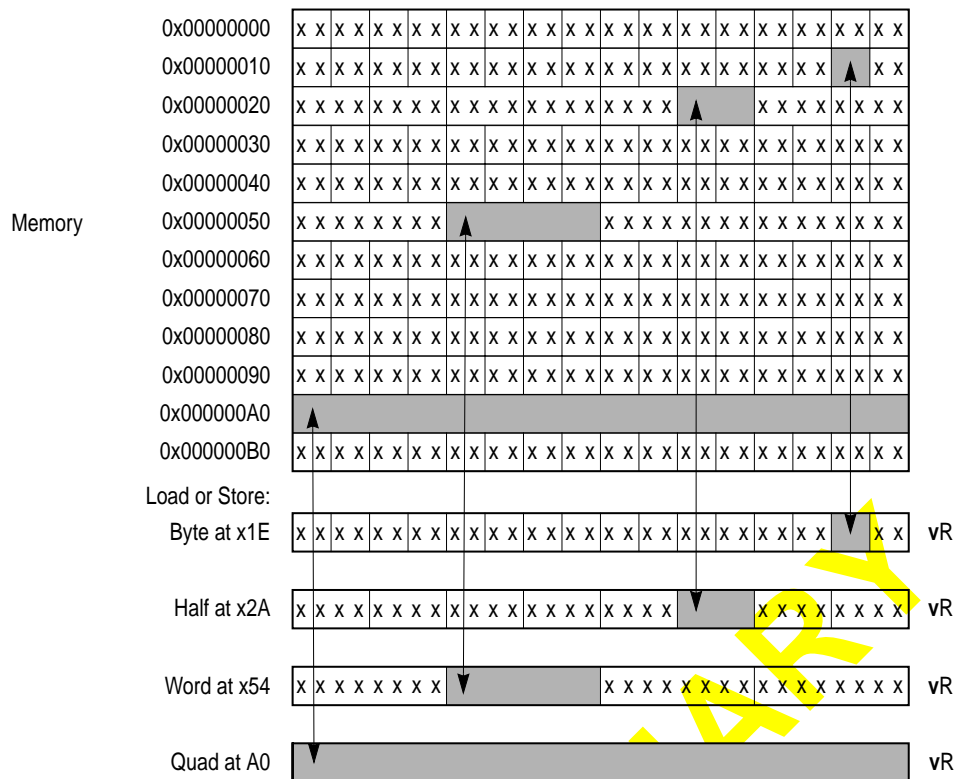
if rA=0 then b ← 0
else      b ← (rA)
EA ← b + (rB)
eb ← EA60-63
vD ← undefined
if the processor is in big-endian mode
then vDeb*8:eb*8+7 ← MEM(EA,1)
else vD120-eb*8:127-eb*8 ← MEM(EA,1)

```

EA = (rA|0)+(rB); m = EA[60–63] (the offset of the byte in its aligned quadword). For big-endian mode, the byte addressed by EA is loaded into byte m of vD. In little-endian mode, it is loaded into byte 15–m of vD. Remaining bytes in vD are undefined.

Other registers altered:

- None



Note: In vector registers, x means boundedly undefined after a load and don't care after a store. In memory, x means don't care after a load, and leave at current value after a store.


**Figure 6-7. Load Instructions**

# Ivehx

Load Vector Element Half Word Indexed

# Ivehx

**Ivehx** **vD,rA,rB**



```

if rA=0 then b ← 0
else      b ← (rA)
EA ← (b + (rB)) & 0xFFFF_FFFF_FFFF_FFFE
eb ← EA60-63
vD ← undefined
if the processor is in big-endian mode
then vDebx8:ebx8+15 ← MEM(EA,2)
else vD112-ebx8:127-ebx8 ← MEM(EA,2)

```

Let the EA be the result of ANDing the sum (rA|0)+(rB) with 0xFFFF\_FFFF\_FFFF\_FFFE. Let m = EA[60–62]; m is the half-word offset of the half-word in its aligned quadword in memory.

If the processor is in big-endian mode, the half-word addressed by EA is loaded into half-word m of vD. If the processor is in little-endian mode, the half-word addressed by EA is loaded into half-word 7–m of vD. The remaining half-words in vD are set to undefined values. Figure 6-7 shows this instruction.

Other registers altered:


- None

# lviewx

Load Vector Element Word Indexed

# lviewx

**lviewx** **vD,rA,rB**



```

if rA=0 then b ← 0
else      b ← (rA)
EA ← (b + (rB)) & 0xFFFF_FFFF_FFFF_FFFC
eb ← EA60-63
vD ← undefined
if the processor is in big-endian mode
then vDeb*8:eb*8+31 ← MEM(EA,4)
else vD96-eb*8:127-eb*8 ← MEM(EA,4)

```

Let the EA be the result of ANDing the sum (rA|0)+(rB) with 0xFFFF\_FFFF\_FFFF\_FFFC. Let m = EA[60–61]; m is the word offset of the word in its aligned quadword in memory.

If the processor is in big-endian mode, the word addressed by EA is loaded into word m of vD. If the processor is in little-endian mode, the word addressed by EA is loaded into word 3–m of vD. The remaining words in vD are set to undefined values. Figure 6-7 shows this instruction.

Other registers altered:

- None

**lvsl** **vD,rA,rB**

31	vD	A	B	6	0
0	5 6	10 11	15 16	20 21	30 31

```


if rA = 0 then b ← 0
else b ← (rA)
addr0-63 ← b + (rB)
sh ← addr60-63
if sh = 0x0 then (vD)0-127 ← 0x000102030405060708090A0B0C0D0E0F
if sh = 0x1 then (vD)0-127 ← 0x0102030405060708090A0B0C0D0E0F10
if sh = 0x2 then (vD)0-127 ← 0x02030405060708090A0B0C0D0E0F1011
if sh = 0x3 then (vD)0-127 ← 0x030405060708090A0B0C0D0E0F101112
if sh = 0x4 then (vD)0-127 ← 0x0405060708090A0B0C0D0E0F10111213
if sh = 0x5 then (vD)0-127 ← 0x05060708090A0B0C0D0E0F1011121314
if sh = 0x6 then (vD)0-127 ← 0x060708090A0B0C0D0E0F101112131415
if sh = 0x7 then (vD)0-127 ← 0x0708090A0B0C0D0E0F10111213141516
if sh = 0x8 then (vD)0-127 ← 0x08090A0B0C0D0E0F1011121314151617
if sh = 0x9 then (vD)0-127 ← 0x090A0B0C0D0E0F101112131415161718
if sh = 0xA then (vD)0-127 ← 0x0A0B0C0D0E0F10111213141516171819
if sh = 0xB then (vD)0-127 ← 0x0B0C0D0E0F101112131415161718191A
if sh = 0xC then (vD)0-127 ← 0x0C0D0E0F101112131415161718191A1B
if sh = 0xD then (vD)0-127 ← 0x0D0E0F101112131415161718191A1B1C
if sh = 0xE then (vD)0-127 ← 0x0E0F101112131415161718191A1B1C1D
if sh = 0xF then (vD)0-127 ← 0x0F101112131415161718191A1B1C1D1E

```

Let the EA be the sum  $(rA|0)+(rB)$ . Let  $sh = EA[60-63]$ . Let  $X$  be the 32-byte value  $0x00 \parallel 0x01 \parallel 0x02 \parallel \dots \parallel 0x1E \parallel 0x1F$ . Bytes  $sh:sh+15$  of  $X$  are placed into  $vD$ . Figure 6-8 shows how this instruction works.


Other registers altered:

- None



**Figure 6-8. Load Vector for Shift Left**

The above **lvsl** instruction followed by a Vector Permute (**vperm**) would do a simulated alignment of a four-element floating-point vector misaligned on quad-word boundary at address  $0x0....C$ .



**Figure 6-9. Instruction vperm Used in Aligning Data**

PRELIMINARY

# lvsr

Load Vector for Shift Right

# lvsr

**lvsr**

**vD,rA,rB**

31	vD	A	B	38	0
0	5 6	10 11	15 16	20 21	30 31

```

if rA = 0 then b ← 0
else          b ← (rA)
EA ← b + (rB)
sh ← EA60-63
if sh=0x0 then vD ← 0x101112131415161718191A1B1C1D1E1F
if sh=0x1 then vD ← 0x0F101112131415161718191A1B1C1D1E
if sh=0x2 then vD ← 0x0E0F101112131415161718191A1B1C1D
if sh=0x3 then vD ← 0x0D0E0F101112131415161718191A1B1C
if sh=0x4 then vD ← 0x0C0D0E0F101112131415161718191A1B
if sh=0x5 then vD ← 0x0B0C0D0E0F101112131415161718191A
if sh=0x6 then vD ← 0x0A0B0C0D0E0F10111213141516171819
if sh=0x7 then vD ← 0x090A0B0C0D0E0F101112131415161718
if sh=0x8 then vD ← 0x08090A0B0C0D0E0F1011121314151617
if sh=0x9 then vD ← 0x0708090A0B0C0D0E0F10111213141516
if sh=0xA then vD ← 0x060708090A0B0C0D0E0F101112131415
if sh=0xB then vD ← 0x05060708090A0B0C0D0E0F1011121314
if sh=0xC then vD ← 0x0405060708090A0B0C0D0E0F10111213
if sh=0xD then vD ← 0x030405060708090A0B0C0D0E0F101112
if sh=0xE then vD ← 0x02030405060708090A0B0C0D0E0F1011
if sh=0xF then vD ← 0x0102030405060708090A0B0C0D0E0F10

```

Let the EA be the sum (rA|0)+(rB). Let sh = EA[60-63]. Let X be the 32-byte value 0x00 || 0x01 || 0x02 || ... || 0x1E || 0x1F. Bytes 16–sh:31–sh of X are placed into vD.

Note that **lvsl** and **lvsr** can be used to create the permute control vector to be used by a subsequent **vperm** instruction. Let X and Y be the contents of vA and vB specified by the **vperm**. The control vector created by **lvsl** causes the **vperm** to select the high-order 16 bytes of the result of shifting the 32-byte value X || Y left by sh bytes. The control vector created by **vsr** causes the **vperm** to select the low-order 16 bytes of the result of shifting X || Y right by sh bytes.


These instructions can also be used to rotate or shift the contents of a **lvsl** or **lvsr** by sh bytes. For rotating, the vector register to be rotated should be specified as both vA and vB for **vperm**. For shifting left, the vB register for **vperm** should contain all zeros and vA should contain the value to be shifted, and vice versa for shifting right.

Other registers altered:

- None



**lvx** **vD,rA,rB** (LRU = 0)



```

if rA=0 then b ← 0
else      b ← (rA)
EA ← (b + (rB)) & 0xFFFF_FFFF_FFFF_FFF0
if the processor is in big-endian mode
then vD ← MEM(EA,16)
else vD ← MEM(EA+8,8) || MEM(EA,8)

```

Let the EA be the result of ANDing the sum (rA|0)+(rB) with 0xFFFF\_FFFF\_FFFF\_FFF0.

If the processor is in big-endian mode, the quadword in memory addressed by EA is loaded into vD.

If the processor is in little-endian mode, the doubleword addressed by EA is loaded into vD[64–127] and the doubleword addressed by EA+8 is loaded into vD[0–63].


Note that on some implementations, the hint provided by the **lvxl** instruction and the corresponding hint provided by the Store Vector Indexed LRU (**stvxl**) instruction (see Section 5.1.1.2, “Prioritizing Cache Block Replacement”) are applied to the entire cache block containing the specified quadword. On such implementations, the effect of the hint may be to cause that cache block to be considered a likely candidate for reuse when space is needed in the cache for a new block. Thus, on such implementations, the hint should be used with caution if the cache block containing the quadword also contains data that may be needed by the program in the near future. Also, the hint may be used before the last reference in a sequence of references to the quadword if the subsequent references are likely to occur sufficiently soon that the cache block containing the quadword is not likely to be displaced from the cache before the last reference. Figure 6-7 shows this instruction.

Other registers altered:

- None

## Load Vector Indexed LRU

**lvxl** **vD,rA,rB** (LRU = 1)



```

if rA=0 then b ← 0
else      b ← (rA)
EA ← (b + (rB)) & 0xFFFF_FFFF_FFFF_FFF0
if the processor is in big-endian mode
    then vD ← MEM(EA,16)
    else vD ← MEM(EA+8,8) || MEM(EA,8)

```

Let the EA be the result of ANDing the sum (rA|0)+(rB) with 0xFFFF\_FFFF\_FFFF\_FFF0.

If the processor is in big-endian mode, the quadword addressed by EA is loaded into vD.

If the processor is in little-endian mode, the doubleword addressed by EA is loaded into vD[64–127] and the doubleword addressed by EA+8 is loaded into vD[0–63].

**lvxl** provides a hint that the quadword addressed by EA will probably not be needed again by the program in the near future.

Note that on some implementations, the hint provided by the **lvxl** instruction and the corresponding hint provided by the Store Vector Indexed LRU (**stvxli**) instruction (see Section 5.1.1.2, “Prioritizing Cache Block Replacement”) are applied to the entire cache block containing the specified quadword. On such implementations, the effect of the hint may be to cause that cache block to be considered a likely candidate for reuse when space is needed in the cache for a new block. Thus, on such implementations, the hint should be used with caution if the cache block containing the quadword also contains data that may be needed by the program in the near future. Also, the hint may be used before the last reference in a sequence of references to the quadword if the subsequent references are likely to occur sufficiently soon that the cache block containing the quadword is not likely to be displaced from the cache before the last reference. Figure 6-7 shows this instruction.

Other registers altered:

- None

# mfvscr

Move from Vector Status and Control Register

# mfvscr

**mfvscr**

**vD**

04	vD	00000	00000	1540	0
0	5 6	10 11	15 16	20 21	30 31

$$vD \leftarrow {}^{96}0 \parallel (VSCR)$$

The contents of the VSCR are placed into vD.

Note that the programmer should assume that **mtvscr** and **mfvscr** will take substantially longer to execute than other VX instructions

Other registers altered:

- None

PRELIMINARY

# mtvscr

Move to Vector Status and Control Register

# mtvscr

**mtvscr**

**vB**

04	0 0 0 0 0	0 0 0 0 0	vB	1604	0
0	5 6	10 11	15 16	20 21	30 31

$$\text{VSCR} \leftarrow (\text{vB})_{96-127}$$

The contents of **vB** are placed into the VSCR.

Other registers altered:

- None

PRELIMINARY

# stvebx

Store Vector Element Byte Indexed

# stvebx

**stvebx**                      **vS,rA,rB**



```
if rA=0 then b ← 0
else      b ← (rA)
EA ← b + (rB)
eb ← EA60-63
if the processor is in big-endian mode
then MEM(EA,1) ← (VS)eb*8:eb*8+7
else MEM(EA,1) ← (VS)120-eb*8:127-eb*8
```

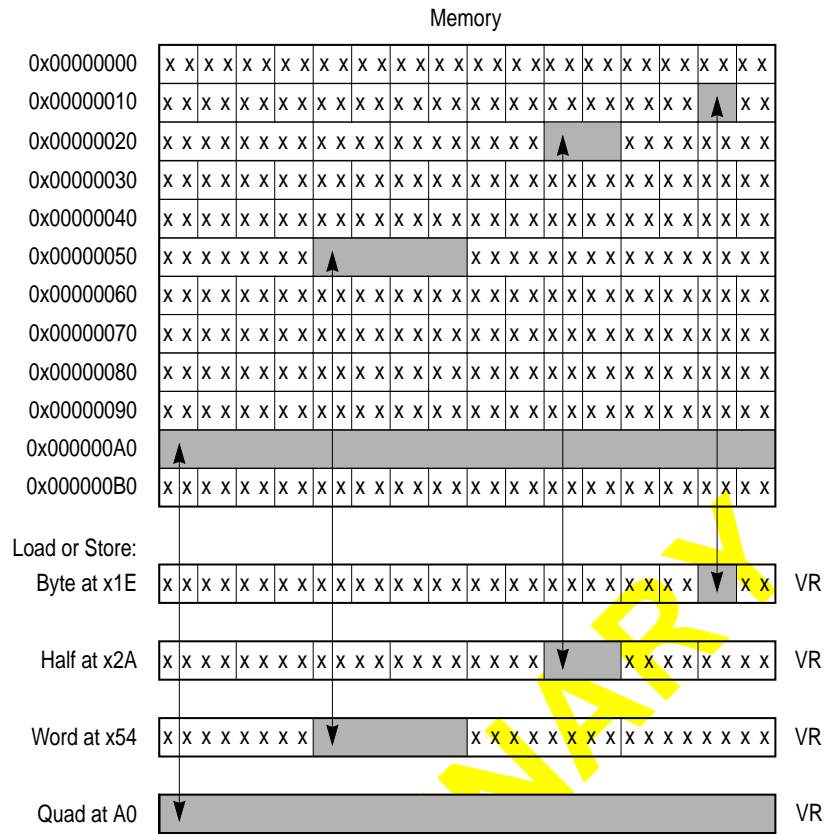
Let the EA be the sum (rA|0)+(rB).

Let m = EA[60–63]; m is the byte offset of the byte in its aligned quadword in memory.

If the processor is in big-endian mode, byte m of VS is stored into the byte in memory addressed by EA. If the processor is in little-endian mode, byte 15–m of VS is stored into the byte addressed by EA.

Other registers altered:

- None



**Note:** In vector registers (VR's), 'x' means 'boundedly undefined' after a load, and 'don't care' after a store. In memory, 'x' means 'don't care' after a load, and 'leave at current value' after a store.


**Figure 6-10. Store Instructions**

# stvehx

Store Vector Element Half Word Indexed

# stvehx

**stvehx**                      **vS,rA,rB**



```

if rA=0 then b ← 0
else      b ← (rA)
EA ← (b + (rB)) & 0xFFFF_FFFF_FFFF_FFFE
eb ← EA60-63
if the processor is in big-endian mode
then MEM(EA,2) ← (VS)eb*8:eb*8+15
else MEM(EA,2) ← (VS)112-eb*8:127-eb*8

```

Let the EA be the result of ANDing the sum (rA|0)+(rB) with 0xFFFF\_FFFF\_FFFF\_FFFE. Let m = EA[60–62]; m is the half-word offset of the half-word in its aligned quadword in memory.

If the processor is in big-endian mode, half-word m of VS is stored into the half-word addressed by EA. If the processor is in little-endian mode, half-word 7–m of VS is stored into the half-word addressed by EA. Figure 6-10 shows of how a store instructions is performed on the vector registers.

Other registers altered:


- None

# stviewx

Store Vector Element Word Indexed

# stviewx

**stviewx**                      **vS,rA,rB**



```
if rA=0 then b ← 0
else      b ← (rA)
EA ← (b + (rB)) & 0xFFFF_FFFF_FFFF_FFFC
eb ← EA60-63
if the processor is in big-endian mode
  then MEM(EA,4) ← (VS)eb*8:eb*8+31
  else MEM(EA,4) ← (VS)96-eb*8:127-eb*8
```

Let the EA be the result of ANDing the sum (rA|0)+(rB) with 0xFFFF\_FFFF\_FFFF\_FFFC. Let m = EA[60-61]; m is the word offset of the word in its aligned quadword in memory.


If the processor is in big-endian mode, word m of VS is stored into the word addressed by EA. If the processor is in little-endian mode, word 3-m of VS is stored into the word addressed by EA. Figure 6-10 shows how a store instructions is performed on the vector registers.

Other registers altered:

- None



**stvx** **vS,rA,rB** (LRU = 0)



```

if rA=0 then b ← 0
else      b ← (rA)
EA ← (b + (rB)) & 0xFFFF_FFFF_FFFF_FFF0
if the processor is in big-endian mode
then MEM(EA,16) ← (VS)
else MEM(EA,16) ← (VS)64-127 || (VS)0-63

```

Let the EA be the result of ANDing the sum (rA|0)+(rB) with 0xFFFF\_FFFF\_FFFF\_FFF0.

If the processor is in big-endian mode, the contents of VS are stored into the quadword addressed by EA. If the processor is in little-endian mode, the contents of VS[64–127] are stored into the doubleword addressed by EA, and the contents of VS[0–63] are stored into the doubleword addressed by EA+8.

**stvxl** and **stvxlt** provide a hint that the quadword addressed by EA will probably not be needed again by the program in the near future

Note that on some implementations, the hint provided by the **lvxl** instruction and the corresponding hint provided by the Store Vector Indexed LRU (**stvxl**) instruction (see Section 5.1.1.2, “Prioritizing Cache Block Replacement”) are applied to the entire cache block containing the specified quadword. On such implementations, the effect of the hint may be to cause that cache block to be considered a likely candidate for reuse when space is needed in the cache for a new block. Thus, on such implementations, the hint should be used with caution if the cache block containing the quadword also contains data that may be needed by the program in the near future. Also, the hint may be used before the last reference in a sequence of references to the quadword if the subsequent references are likely to occur sufficiently soon that the cache block containing the quadword is not likely to be displaced from the cache before the last reference.

Figure 6-10 shows how a store instructions is performed on the vector registers.

Other registers altered:


- None

# stvx1

Store Vector Indexed LRU

# stvx1

**stvx1**                      **vS,rA,rB**                      (LRU = 1)



```

if rA=0 then b ← 0
else      b ← (rA)
EA ← (b + (rB)) & 0xFFFF_FFFF_FFFF_FFF0
if the processor is in big-endian mode
then MEM(EA,16) ← (VS)
else MEM(EA,16) ← (VS)64-127 || (VS)0-63

```

Let the EA be the result of ANDing the sum (rA|0)+(rB) with 0xFFFF\_FFFF\_FFFF\_FFF0.

If the processor is in big-endian mode, the contents of VS are stored into the quadword addressed by EA. If the processor is in little-endian mode, the contents of VS[64–127] are stored into the doubleword addressed by EA, and the contents of VS[0–63] are stored into the doubleword addressed by EA+8.

**stvx1** and **stvx1t** provide a hint that the quad word addressed by EA will probably not be needed again by the program in the near future

Note that on some implementations, the hint provided by the **lvxl** instruction and the corresponding hint provided by the Store Vector Indexed LRU (**stvx1**) instruction (see Section 5.1.1.2, “Prioritizing Cache Block Replacement”) are applied to the entire cache block containing the specified quadword. On such implementations, the effect of the hint may be to cause that cache block to be considered a likely candidate for reuse when space is needed in the cache for a new block. Thus, on such implementations, the hint should be used with caution if the cache block containing the quadword also contains data that may be needed by the program in the near future. Also, the hint may be used before the last reference in a sequence of references to the quadword if the subsequent references are likely to occur sufficiently soon that the cache block containing the quadword is not likely to be displaced from the cache before the last reference.

Figure 6-10 shows how a store instruction is performed on the vector registers.

Other registers altered:

- None

**PRELIMINARY**

# vaddcuw

Vector Add Carryout Unsigned Word

# vaddcuw

**vaddcuw**                      **vD,vA,vB**

04					vD					vA					vB					384												
0					5					10					15					20												
					6					11					16					21												

```
do i=0 to 127 by 32
  aop0-32 ← ZeroExtend((vA)i:i+31,33)
  bop0-32 ← ZeroExtend((vB)i:i+31,33)
  temp0-32 ← aop0-32 +int bop0-32
  vDi:i+31 ← ZeroExtend(temp0,32)
```

Each unsigned-integer word element in **vA** is added to the corresponding unsigned-integer word element in **vB**. The carry out of bit 0 of the 32-bit sum is zero-extended to 32 bits and placed into the corresponding word element of **vD**.

Other registers altered:


- None

# vaddfp

Vector Add Floating Point

# vaddfp

**vaddfp** **vD,vA,vB**




```
do i = 0,127,32
  (vD)i:i+31 ← RndToNearFP32((vA)i:i+31 +fp (vB)i:i+31)
end
```

The four 32-bit floating-point values in **vA** are added to the four 32-bit floating-point values in **vB**. The four intermediate results are rounded and placed in **VD**.

Other registers altered:

- None



**Figure 6-11. Basic 2-Source Operands —32-Bit Elements**

# vaddsbs

Vector Add Signed Byte Saturate

# vaddsbs

**vaddsbs** **vD,vA,vB**

04	vD	vA	vB	768
0	5 6	10 11	15 16	20 21
				31

```

n ← LENGTH(element)
do i=0 to 127 by n
  aop0:n ← SignExtend((vA)i:i+n-1,n+1)
  bop0:n ← SignExtend((vB)i:i+n-1,n+1)
  temp0:n ← aop0:n +int bop0:n
  vDi:i+n-1 ← SItoSIsat(temp0:n,n)

```

For **vaddsbs** each element is a byte.

Each signed-integer element in **vA** is added to the corresponding signed-integer element in **vB**.


If the intermediate result is greater than  $2^{n-1}-1$  it saturates to  $2^{n-1}-1$  and if it is less than  $-2^{n-1}$  it saturates to  $-2^{n-1}$ , where n is the length of the element.

The signed-integer result is placed into the corresponding element of **VD**.

Other registers altered:

- Vector status and control register (VSCR) Vector status and control register (VSCR):

Affected: SAT



**Figure 6-12. Basic 2-Source Operands, Sixteen 8-Bit Elements**

# vaddshs

Vector Add Signed Half Word Saturate

# vaddshs

**vaddshs** **vD,vA,vB**

04	vD	vA	vB	832
0	5 6	10 11	15 16	20 21
				31

```

n ← LENGTH(element)
do i=0 to 127 by n
  aop0:n ← SignExtend((vA)i:i+n-1,n+1)
  bop0:n ← SignExtend((vB)i:i+n-1,n+1)
  temp0:n ← aop0:n +int bop0:n
  vDi:i+n-1 ← SToSIsat(temp0:n,n)

```

For **vaddshs** each element is a half word .

Each signed-integer element in **vA** is added to the corresponding signed-integer element in **vB**.


If the intermediate result is greater than  $2^{n-1}-1$  it saturates to  $2^{n-1}-1$  and if it is less than  $-2^{n-1}$  it saturates to  $-2^{n-1}$ , where n is the length of the element.

The signed-integer result is placed into the corresponding element of **vD**.

Other registers altered:

- Vector status and control register (VSCR):

Affected: SAT



**Figure 6-13. Basic 2-Source Operands, Eight 16-Bit Elements**

**PRELIMINARY**



# vaddsws

Vector Add Signed Word Saturate

# vaddsws

**vaddsws** **vD,vA,vB**

04	vD	vA	vB	896
0	5 6	10 11	15 16	20 21
				31

```

n ← LENGTH(element)
do i=0 to 127 by n
  aop0:n ← SignExtend((vA)i:i+n-1,n+1)
  bop0:n ← SignExtend((vB)i:i+n-1,n+1)
  temp0:n ← aop0:n +int bop0:n
  vDi:i+n-1 ← SItoSIsat(temp0:n,n)

```

For **vaddsws** each element is a word.

Each signed-integer element in **vA** is added to the corresponding signed-integer element in **vB**.


If the intermediate result is greater than  $2^{n-1}-1$  it saturates to  $2^{n-1}-1$  and if it is less than  $-2^{n-1}$  it saturates to  $-2^{n-1}$ , where  $n$  is the length of the element.

The signed-integer result is placed into the corresponding element of **vD**.

Other registers altered:

- Vector status and control register (VSCR):

Affected: SAT



**Figure 6-14. Basic 2-Source Operands, Four 32-bit Elements**

# vaddubm

Vector Add Unsigned Byte Modulo

# vaddubm

**vaddubm** **vD,vA,vB**

04	vD	vA	vB	0
0	5 6	10 11	15 16	20 21
				31

```

n ← LENGTH(element)
do i=0 to 127 by n
    vDi:i+n-1 ← (vA)i:i+n +int (vB)i:i+n-1

```


For **vaddubm** each element is a byte.

Each unsigned-integer element in **vA** is added to the corresponding unsigned-integer element in **vB**. The unsigned-integer result is placed into the corresponding element of **vD**.

Note that the **vaddubm** instruction can be used for unsigned or signed integers.

Other registers altered:

- None



**Figure 6-15. Basic 2-Source Operands, Sixteen 8-bit Elements**

# vaddubs

Vector Add Unsigned Byte Saturate

# vaddubs

**vaddubs** **vD,vA,vB**

04	vD	vA	vB	512
0	5 6	10 11	15 16	20 21
				31

```

n ← LENGTH(element)
do i=0 to 127 by n
  aop0:n ← ZeroExtend((vA)i:i+n-1,n+1)
  bop0:n ← ZeroExtend((vB)i:i+n-1,n+1)
  temp0:n ← aop0:n +int bop0:n
  vDi:i+n-1 ← UItoUISat(temp0:n,n)

```

For **vaddubs** each element is a byte.

Each unsigned-integer element in **vA** is added to the corresponding unsigned-integer element in **vB**.


If the intermediate result is greater than  $2^n-1$  it saturates to  $2^n-1$ , where  $n$  is the length of the element.

The unsigned-integer result is placed into the corresponding element of **vD**.

Other registers altered:

- Vector status and control register (VSCR):

Affected: SAT



**Figure 6-16. Basic 2-Source Operands, Sixteen 8-bit Elements**

# vadduhm

Vector Add Unsigned Half Word Modulo

# vadduhm

**vadduhm** **vD,vA,vB**

04	vD	vA	vB	64
0	5 6	10 11	15 16	20 21
				31

```

n ← LENGTH(element)
do i=0 to 127 by n
    vDi:i+n-1 ← (vA)i:i+n +int (vB)i:i+n-1

```

For **vadduhm** each element is a half word.

Each unsigned-integer element in **vA** is added to the corresponding unsigned-integer element in **vB**. The unsigned-integer result is placed into the corresponding element of **vD**.

Note that the **vadduhm** instruction can be used for unsigned or signed integers.

Other registers altered:

- None




Figure 6-17. Basic 2-Source Operands, Eight 16-bit Elements

# vadduhs

Vector Add Unsigned Half Word Saturate

# vadduhs

**vadduhs** **vD,vA,vB**

04	vD	vA	vB	576
0	5 6	10 11	15 16	20 21
				31

```

n ← LENGTH(element)
do i=0 to 127 by n
  aop0:n ← ZeroExtend((vA)i:i+n-1,n+1)
  bop0:n ← ZeroExtend((vB)i:i+n-1,n+1)
  temp0:n ← aop0:n +int bop0:n
  vDi:i+n-1 ← UItoUISat(temp0:n,n)

```

For **vadduhs** each element is a half word.

Each unsigned-integer element in **vA** is added to the corresponding unsigned-integer element in **vB**.

If the intermediate result is greater than  $2^n-1$  it saturates to  $2^n-1$ , where n is the length of the element.

The unsigned-integer result is placed into the corresponding element of **vD**.

Other registers altered:

- Vector status and control register (VSCR):

Affected: SAT




Figure 6-18. Basic 2-Source Operands, Eight 16-bit Elements

**vadduwm**


**vadduwm** **vD,vA,vB**

For **vadduwm** each element is a word.

Note that the **vadduwm** instruction can be used for unsigned or signed integers.

Other registers altered:

- None



**Figure 6-19. Basic 2-Source Operands, Four 32-bit Elements**

# vadduws

Vector Add Unsigned Word Saturate

# vadduws

**vadduws** **vD,vA,vB**

04	vD	vA	vB	640
0	5 6	10 11	15 16	20 21
				31

```

n ← LENGTH(element)
do i=0 to 127 by n
  aop0:n ← ZeroExtend((vA)i:i+n-1,n+1)
  bop0:n ← ZeroExtend((vB)i:i+n-1,n+1)
  temp0:n ← aop0:n +int bop0:n
  vDi:i+n-1 ← UItoUISat(temp0:n,n)

```

For **vadduws** each element is a word.

Each unsigned-integer element in **vA** is added to the corresponding unsigned-integer element in **vB**.

If the intermediate result is greater than  $2^n-1$  it saturates to  $2^n-1$ , where n is the length of the element.

The unsigned-integer result is placed into the corresponding element of **vD**.

Other registers altered:

- Vector status and control register (VSCR):

Affected: SAT





Figure 6-20. Basic 2-Source Operands, Four 32-bit Elements

# vand

Vector Logical AND

# vand

**vand** **vD,vA,vB**



$vD \leftarrow (vA) \& (vB)$   
The contents of **vA** are ANDed with the contents of **vB** and the result is placed into **vD**.

- Other registers altered:
- None




Figure 6-21. Basic 2-Source Operands




# vandc

Vector Logical AND with Complement

# vandc

**vand** **vD,vA,vB**



$vD \leftarrow (vA) \& \neg(vB)$   
The contents of **vA** are ANDed with the complement of the contents of **vB** and the result is placed into **vD**.

- Other registers altered:
- None




Figure 6-22. Basic 2-Source Operands

# vavg**sb**

Vector Average Signed Byte

# vavg**sb**

**vavg**sb****                      **vD,vA,vB**

04	<b>vD</b>	<b>vA</b>	<b>vB</b>	1282
0	5 6	10 11	15 16	20 21
				31

```

n ← LENGTH(element)
do i=0 to 127 by n
  aop0:n ← SignExtend((vA)i:i+n-1,n+1)
  bop0:n ← SignExtend((vB)i:i+n-1,n+1)
  temp0:n ← aop0:n +int bop0:n +int 1
  vDi:i+n-1 ← temp0:n-1


```

For **vavg**sb**** each element is a byte.

Each signed-integer element in **vA** is added to the corresponding signed-integer element in **vB**, producing an (n+1)-bit signed-integer sum where **n** is the length of the element. The sum is incremented by 1. The high-order **n** bits(0–7, 8–15, etc...) of the result are placed into the corresponding element of **vD**.

Other registers altered:

- None



**Figure 6-23. Basic 2-Source Operands—Sixteen 8-bit Elements**

# vavgsh

Vector Average Signed Half Word

# vavgsh

**vavgsh** **vD,vA,vB**

04	vD	vA	vB	1346
0	5 6	10 11	15 16	20 21
				31

```

n ← LENGTH(element)
do i=0 to 127 by n
  aop0:n ← SignExtend((vA)i:i+n-1,n+1)
  bop0:n ← SignExtend((vB)i:i+n-1,n+1)
  temp0:n ← aop0:n +int bop0:n +int 1
  vDi:i+n-1 ← temp0:n-1


```

For **vavgsh** each element is a half word.

Each signed-integer element in **vA** is added to the corresponding signed-integer element in **vB**, producing an (n+1)-bit signed-integer sum where **n** is the length of the element. The sum is incremented by 1. The high-order **n** bits (0-15,16-31, etc...) of the result are placed into the corresponding element of **vD**.

Other registers altered:

- None



**Figure 6-24. Basic 2-Source Operands, Eight 16-bit Elements**

# vavgsw

Vector Average Signed Word

# vavgsw

**vavgsw** **vD,vA,vB**

04	vD	vA	vB	1410
0	5 6	10 11	15 16	20 21
				31

```

n ← LENGTH(element)
do i=0 to 127 by n
  aop0:n ← SignExtend((vA)i:i+n-1,n+1)
  bop0:n ← SignExtend((vB)i:i+n-1,n+1)
  temp0:n ← aop0:n +int bop0:n +int 1
  vDi:i+n-1 ← temp0:n-1


```

For **vavgsw** each element is a word.

Each signed-integer element in **vA** is added to the corresponding signed-integer element in **vB**, producing an (n+1)-bit signed-integer sum where **n** is the length of the element. The sum is incremented by 1. The high-order **n** bits(0–31, 32–63, etc...) of the result are placed into the corresponding element of **vD**.

Other registers altered:

- None



**Figure 6-25. Basic 2-Source Operands, Four 32-bit Elements**

# vavgub

Vector Average Unsigned Byte

# vavgub

**vavgub** **vD,vA,vB**

04	vD	vA	vB	1026
0	5 6	10 11	15 16	20 21
				31

```

n ← LENGTH(element)
do i=0 to 127 by n
  aop0:n ← ZeroExtend((vA)i:i+n-1,n+1)
  bop0:n ← ZeroExtend((vB)i:i+n-1,n+1)
  temp0:n ← aop0:n +int bop0:n +int 1
  vDi:i+n-1 ← temp0:n-1


```

For **vavgub** each element is a byte.

Each unsigned-integer element in **vA** is added to the corresponding unsigned-integer element in **vB**, producing an (n+1)-bit unsigned-integer sum where n is the length of the element. The sum is incremented by 1. The high-order n bits (0–7, 8–15, etc...) of the result are placed into the corresponding element of **vD**.

Other registers altered:

- None



**Figure 6-26. Basic 2-Source Operands—Sixteen 8-bit Elements**

# vavguh

Vector Average Unsigned Half Word

# vavguh

**vavguh** **vD,vA,vB**

04	vD	vA	vB	1090
0	5 6	10 11	15 16	20 21
				31

```

n ← LENGTH(element)
do i=0 to 127 by n
  aop0:n ← ZeroExtend((vA)i:i+n-1,n+1)
  bop0:n ← ZeroExtend((vB)i:i+n-1,n+1)
  temp0:n ← aop0:n +int bop0:n +int 1
  vDi:i+n-1 ← temp0:n-1


```

For **vavguh** each element is a half word.

Each unsigned-integer element in **vA** is added to the corresponding unsigned-integer element in **vB**, producing an (n+1)-bit unsigned-integer sum where n is the length of the element. The sum is incremented by 1. The high-order n bits (0-15,16-31, etc...) of the result are placed into the corresponding element of **vD**.

Other registers altered:

- None



**Figure 6-27. Basic 2-Source Operands, Eight 16-bit Elements**

# vavguw

Vector Average Unsigned Word

# vavguw

**vavguw** **vD,vA,vB**

04	vD	vA	vB	1154
0	5 6	10 11	15 16	20 21
				31

```

n ← LENGTH(element)
do i=0 to 127 by n
  aop0:n ← ZeroExtend((vA)i:i+n-1,n+1)
  bop0:n ← ZeroExtend((vB)i:i+n-1,n+1)
  temp0:n ← aop0:n +int bop0:n +int 1
  vDi:i+n-1 ← temp0:n-1


```

For **vavguw** each element is a word.

Each unsigned-integer element in **vA** is added to the corresponding unsigned-integer element in **vB**, producing an (n+1)-bit unsigned-integer sum where n is the length of the element. The sum is incremented by 1. The high-order n bits (0–31, 32–63, etc...) of the result are placed into the corresponding element of **vD**.

Other registers altered:

- None



**Figure 6-28. Basic 2-Source Operands, Four 32-bit Elements**

# vcfsx

Vector Convert from Signed Fixed-Point Word

# vcfsx

**vcfsx** **vD,vB,UIMM**

04	vD	UIMM	vB	842
0	5 6	10 11	15 16	20 21
				31


do i=0 to 127 by 32

$vD_{i:i+31} \leftarrow \text{CnvtSI32ToFP32}((vB)_{i:i+31}) \div_{fp} 2^{UIMM}$

Each signed fixed-point integer word element in **vB** (word = 32 bits, 4 total elements) is converted to the nearest single-precision floating-point value. The result is divided by  $2^{UIMM}$  and placed into the corresponding word element of **vD**.

Other registers altered:

- None



**Figure 6-29. Convert Four 32-bit Integer Elements to Four 32-bit Floating Point Elements**



Vector Convert from Unsigned Fixed-Point Word

vcfux                      vD,vB,UIMM

04	vD	UIMM	vB	778
0	5 6	10 11	15 16	20 21
				31


do i=0 to 127 by 32

$$vD_{i:i+31} \leftarrow \text{CnvtUI32ToFP32}((vB)_{i:i+31}) \div_{fp} 2^{UIMM}$$

Each unsigned fixed-point integer word element in **vB** (word = 32 bits, 4 total elements) is converted to the nearest single-precision floating-point value. The result is divided by  $2^{UIMM}$  and placed into the corresponding word element of **vD**.

Other registers altered:

- None



**Figure 6-30. Convert Four 32-bit Integer Elements to Four 32-bit Floating Point Elements**

# vcmpbfp

Vector Compare Bounds Floating Point

# vcmpbfp

**vcmpbfp** **vD,vA,vB** (Rc = 0)  
**vcmpbfp.** **vD,vA,vB** (Rc = 1)

04	vD	vA	vB	Rc	966
0	5 6	10 11	15 16	20 21 22	31

```

do i=0 to 127 by 32
  le ← ((vA)i:i+31 ≤fp (vB)i:i+31)
  ge ← ((vA)i:i+31 ≥fp -(vB)i:i+31)
  vDi:i+31 ← ¬le || ¬ge || 300
if Rc=1 then do
  ib ← (vD = 1280)
  CR24-27 ← 0b00 || ib || 0b0

```

Each single-precision word element in **vA** (word = 32 bits, 4 total elements) is compared to the corresponding element in **vB**. A 2-bit value is formed that indicates whether the element in **vA** is within the bounds specified by the element in **vB**, as follows.

Bit 0 of the 2-bit value is zero if the element in **vA** is less than or equal to the element in **vB**, and is one otherwise. Bit 1 of the 2-bit value is zero if the element in **vA** is greater than or equal to the negation of the element in **vB**, and is one otherwise.

The 2-bit value is placed into the high-order two bits of the corresponding word element (high order bits 0–1 for word element 1, high order bits 32–33 for word element 2, high order bits 64–65 for word element 3, high order bits 96–97 for word element 4) of **vD** and the remaining bits of the element are cleared.


If Rc=1, CR Field 6 is set to indicate whether all four elements in **vA** are within the bounds specified by the corresponding element in **vB**, as follows.

CR6 = 0b00 || all\_within\_bounds || 0

Note that each single-precision floating-point word element in **vB** should be non-negative; if it is negative, the corresponding element in **vA** will necessarily be out of bounds.

Other registers altered:

- Condition register (CR6 field):  
Affected: Bit 2(if Rc = 1)



**Figure 6-31. Basic 2-Source Operands , 32-Bit Elements**

PRELIMINARY

# vcmpeqfp x

Vector Compare Equal-to-Floating Point

# vcmpeqfp x

**vcmpeqfp**                      **vD,vA,vB**  
**vcmpeqfp.**                    **vD,vA,vB**

04	vD	vA	vB	Rc	198
0	5 6	10 11	15 16	20 21 22	31

```

do i=0 to 127 by 32
  if (vA)i:i+31 =fp (vB)i:i+31
    then vDi:i+31 ← 0xFFFF_FFFF
    else vDi:i+31 ← 0x0000_0000
if Rc=1 then do
  t ← (vD = 1281)
  f ← (vD = 1280)
  CR24-27 ← t || 0b0 || f || 0b0


```

Each single-precision floating-point word element in **vA** (word = 32 bits, 4 total elements) is compared to the corresponding single-precision floating-point word element in **vB**. The corresponding word element in **vD** is set to all 1s if the element in **vA** is equal to the element in **vB**, and is cleared to all 0s otherwise.

Other registers altered:

- Condition register (CR6 field):

Affected: Bits 0-3(if Rc = 1)



**Figure 6-32. Basic 2-Source Operands, Four 32-bit Elements**

# vcmpequbx

Vector Compare Equal-to Unsigned Byte

# vcmpequbx

**vcmpequb** **vD,vA,vB**

**vcmpequb.** **vD,vA,vB**

04	vD	vA	vB	Rc	6
0	5 6	10 11	15 16	20 21 22	31

```

n ← LENGTH(element)
do i=0 to 127 by n
  if (vA)i:i+n-1 =int (vB)i:i+n-1
    then vDi:i+n-1 ← n1
    else vDi:i+n-1 ← n0
if Rc=1 then do
  t ← (vD = 1281)
  f ← (vD = 1280)
  CR[24-27] ← t || 0b0 || f || 0b0

```

For **vcmpequb** each element is a byte.


Each unsigned-integer element in **vA** is compared to the corresponding unsigned-integer element in **vB**. The corresponding element in **vD** is set to all 1s if the element in **vA** is equal to the element in **vB**, and is cleared to all 0s otherwise.

Note that **vcmpequb[.]** can be used for unsigned or signed integers.

Other registers altered:

- Condition register (CR6 field):

Affected: Bits 0-3 (if Rc = 1)



**Figure 6-33. Basic 2-Source Operands—Sixteen 8-bit Elements**

**PRELIMINARY**

# vcmpequhx

Vector Compare Equal-to Unsigned Half Word

# vcmpequhx

**vcmpequh** **vD,vA,vB**

**vcmpequh.** **vD,vA,vB**

04	vD	vA	vB	Rc	70
0	5 6	10 11	15 16	20 21 22	31

```

n ← LENGTH(element)
do i=0 to 127 by n
  if (vA)i:i+n-1 =int (vB)i:i+n-1
    then vDi:i+n-1 ← n1
    else vDi:i+n-1 ← n0
if Rc=1 then do
  t ← (vD = 1281)
  f ← (vD = 1280)
  CR[24-27] ← t || 0b0 || f || 0b0

```

For **vcmpequh** each element is a half word.

Each unsigned-integer element in **vA** is compared to the corresponding unsigned-integer element in **vB**. The corresponding element in **vD** is set to all 1s if the element in **vA** is equal to the element in **vB**, and is cleared to all 0s otherwise.

Note that **vcmpequh[.]** can be used for unsigned or signed integers.

Other registers altered:

- Condition register (CR6 field):

Affected: Bits 0–3(if Rc = 1)




Figure 6-34. Basic 2-Source Operands, Eight 16-bit Elements

# vcmpequwx

Vector Compare Equal-to Unsigned Word

# vcmpequwx

**vcmpequw** **vD,vA,vB**

**vcmpequw.** **vD,vA,vB**

04	vD	vA	vB	Rc	134
0	5 6	10 11	15 16	20 21 22	31

```

n ← LENGTH(element)
do i=0 to 127 by n
  if (vA)i:i+n-1 =int (vB)i:i+n-1
    then vDi:i+n-1 ← n1
    else vDi:i+n-1 ← n0
if Rc=1 then do
  t ← (vD = 1281)
  f ← (vD = 1280)
  CR[24-27] ← t || 0b0 || f || 0b0

```

For **vcmpequw** each element is a word.


Each unsigned-integer element in **vA** is compared to the corresponding unsigned-integer element in **vB**. The corresponding element in **vD** is set to all 1s if the element in **vA** is equal to the element in **vB**, and is cleared to all 0s otherwise.

Note that **vcmpequw[.]** can be used for unsigned or signed integers.

Other registers altered:

- Condition register (CR6 field):

Affected: Bits 0-3(if Rc = 1)



**Figure 6-35. Basic 2-Source Operands, Four 32-bit Elements**



**PRELIMINARY**

# vcmpgfp<sub>x</sub>

Vector Compare Greater-than-or-Equal-to Floating Point

# vcmpgfp<sub>x</sub>

**vcmpgfp** **vD,vA,vB** (Rc = 0)

**vcmpgfp.** **vD,vA,vB** (Rc = 1)

04	vD	vA	vB	Rc	454
0	5 6	10 11	15 16	20 21 22	31

```

do i=0 to 127 by 32
  if (vA)i:i+31 ≥fp (vB)i:i+31
    then vDi:i+31 ← 0xFFFF_FFFF
    else vDi:i+31 ← 0x0000_0000
if Rc=1 then do
  t ← (vD = 1281)
  f ← (vD = 1280)
  CR24-27 ← t || 0b0 || f || 0b0


```

Each single-precision floating-point word element in **vA** (word = 32 bits, 4 total elements) is compared to the corresponding single-precision floating-point word element in **vB**. The corresponding word element in **vD** is set to all 1s if the element in **vA** is greater than or equal to the element in **vB**, and is cleared to all 0s otherwise.

Other registers altered:

- Condition register (CR6 field):

Affected: Bits 0-3(if Rc = 1)



**Figure 6-36. Basic 2-Source Operands, Four 32-bit Elements**

# vcmpgtfpx

Vector Compare Greater than Floating-Point

# vcmpgtfpx

**vcmpgtfp**                      **vD,vA,vB**  
**vcmpgtfp.**                    **vD,vA,vB**

04	vD	vA	vB	Rc	710
0	5 6	10 11	15 16	20 21 22	31

```

do i=0 to 127 by 32
  if (vA)i:i+31 >fp (vB)i:i+31
    then vDi:i+31 ← 0xFFFF_FFFF
    else vDi:i+31 ← 0x0000_0000
if Rc=1 then do
  t ← (vD = 1281)
  f ← (vD = 1280)
  CR[24-27] ← t || 0b0 || f || 0b0


```

Each single-precision floating-point word element in **vA** (word = 32 bits, 4 total elements) is compared to the corresponding single-precision floating-point word element in **vB**. The corresponding word element in **vD** is set to all 1s if the element in **vA** is greater than the element in **vB**, and is cleared to all 0s otherwise.

Other registers altered:

- Condition register (CR6 field):

Affected: Bits 0-3(if Rc = 1)



**Basic 2-Source Operands, Four 32-bit Elements**

# vcmpgtsbx

Vector Compare Greater than Signed Byte

# vcmpgtsbx

**vcmpgtsb** **vD,vA,vB**

**vcmpgtsb.** **vD,vA,vB**

04	vD	vA	vB	Rc	774
0	5 6	10 11	15 16	20 21 22	31

```

n ← LENGTH(element)
do i=0 to 127 by n
  if (vA)i:i+n-1 >si (vB)i:i+n-1
    then vDi:i+n-1 ← n1
    else vDi:i+n-1 ← n0
if Rc=1 then do
  t ← (vD = 1281)
  f ← (vD = 1280)
  CR24-27 ← t || 0b0 || f || 0b0

```


For **vcmpgtsb** each element is a byte.

Each signed-integer element in **vA** is compared to the corresponding signed-integer element in **vB**. The corresponding element in **vD** is set to all 1s if the element in **vA** is greater than the element in **vB**, and is cleared to all 0s otherwise.

Other registers altered:

- Condition register (CR6 field):

Affected: Bits 0-3(if Rc = 1)



**Figure 6-37. Basic 2-Source Operands—Sixteen 8-bit Elements**

**PRELIMINARY**

# vcmpgtshx

# vcmpgtshx

Vector Compare Greater thanCondition register Signed Half Word

**vcmpgtsh**  $\mathbf{vD}, \mathbf{vA}, \mathbf{vB}$

**vcmpgtsh.**  $\mathbf{vD}, \mathbf{vA}, \mathbf{vB}$

04	<b>vD</b>	<b>vA</b>	<b>vB</b>	Rc	838
0	5 6	10 11	15 16	20 21 22	31

For **vcmpgtsh** each element is a half word.

Each signed-integer element in **vA** is compared to the corresponding signed-integer element in **vB**. The corresponding element in **vD** is set to all 1s if the element in **vA** is greater than the element in **vB**, and is cleared to all 0s otherwise.

Other registers altered:

- Condition register (CR6 field):

Affected: Bits 0-3(if Rc = 1)




Figure 6-38. Basic 2-Source Operands, Eight 16-bit Elements


# vcmpgtswx

Vector Compare Greater than Signed Word

# vcmpgtswx

**vcmpgtsw**                      **vD,vA,vB**

**vcmpgtsw.**                    **vD,vA,vB**



For **vcmpgtsw** each element is a word.

Each signed-integer element in **vA** is compared to the corresponding signed-integer element in **vB**. The corresponding element in **vD** is set to all 1s if the element in **vA** is greater than the element in **vB**, and is cleared to all 0s otherwise.

Other registers altered:

- Condition register (CR6 field):

Affected: Bits 0-3(if Rc = 1)




Figure 6-39. Basic 2-Source Operands, Four 32-bit Elements

# vcmpgtubx

Vector Compare Greater than Unsigned Byte

# vcmpgtubx

**vcmpgtub**  $\mathbf{vD}, \mathbf{vA}, \mathbf{vB}$

**vcmpgtub.**  $\mathbf{vD}, \mathbf{vA}, \mathbf{vB}$

04	<b>vD</b>	<b>vA</b>	<b>vB</b>	Rc	518
0	5 6	10 11	15 16	20 21 22	31

```

n ← LENGTH(element)
do i=0 to 127 by n
  if (vA)i:i+n-1 >ui (vB)i:i+n-1
    then vDi:i+n-1 ← n1
    else vDi:i+n-1 ← n0
if Rc=1 then do
  t ← (vD = 1281)
  f ← (vD = 1280)
  CR[24-27] ← t || 0b0 || f || 0b0

```


For **vcmpgtub** each element is a byte.

Each unsigned-integer element in **vA** is compared to the corresponding unsigned-integer element in **vB**. The corresponding element in **vD** is set to all 1s if the element in **vA** is greater than the element in **vB**, and is cleared to all 0s otherwise.

Other registers altered:

- Condition register (CR6 field):

Affected: Bits 0-3(if Rc = 1)



**Figure 6-40. Basic 2-Source Operands—Sixteen 8-bit Elements**



**PRELIMINARY**

# vcmpgtuhx

Vector Compare Greater than Unsigned Half Word

# vcmpgtuhx

**vcmpgtuh** **vD,vA,vB**

**vcmpgtuh.** **vD,vA,vB**

04	vD	vA	vB	Rc	582
0	5 6	10 11	15 16	20 21 22	31

```

n ← LENGTH(element)
do i=0 to 127 by n
  if (vA)i:i+n-1 >ui (vB)i:i+n-1
    then vDi:i+n-1 ← n1
    else vDi:i+n-1 ← n0
if Rc=1 then do
  t ← (vD = 1281)
  f ← (vD = 1280)
  CR[24-27] ← t || 0b0 || f || 0b0

```

For **vcmpgtuh** each element is a half word.

Each unsigned-integer element in **vA** is compared to the corresponding unsigned-integer element in **vB**. The corresponding element in **vD** is set to all 1s if the element in **vA** is greater than the element in **vB**, and is cleared to all 0s otherwise.

Other registers altered:

- Condition register (CR6 field):

Affected: Bits 0-3(if Rc = 1)




Figure 6-41. Basic 2-Source Operands, Eight 16-bit Elements

# vcmpgtuw

Vector Compare Greater than Unsigned Word

# vcmpgtuw

**vcmpgtuw** **vD,vA,vB**

**vcmpgtuw.** **vD,vA,vB**

04	vD	vA	vB	Rc	646
0	5 6	10 11	15 16	20 21 22	31

```

n ← LENGTH(element)
do i=0 to 127 by n
  if (vA)i:i+n-1 >ui (vB)i:i+n-1
    then vDi:i+n-1 ← n1
    else vDi:i+n-1 ← n0
if Rc=1 then do
  t ← (vD = 1281)
  f ← (vD = 1280)
  CR[24-27] ← t || 0b0 || f || 0b0

```


For **vcmpgtuw** each element is a word.

Each unsigned-integer element in **vA** is compared to the corresponding unsigned-integer element in **vB**. The corresponding element in **vD** is set to all 1s if the element in **vA** is greater than the element in **vB**, and is cleared to all 0s otherwise.

Other registers altered:

- Condition register (CR6 field):

Affected: Bits 0-3(if Rc = 1)



**Figure 6-42. Basic 2-Source Operands, Four 32-bit Elements**

# vctxsxs

# vctxsxs

Vector Convert to Signed Fixed-Point Word Saturate

**vctxsxs**                      **vD,vB,UIMM**

04	vD	UIMM	vB	970
0	5 6	10 11	15 16	20 21
				31

```

do i=0 to 127 by 32
  if (vB)i+1:i+8=255 | (vB)i+1:i+8 + UIMM ≤ 254 then
    vDi:i+31 ← CnvtFP32ToSI32Sat((vB)i:i+31 ×fp 2UIMM)
  else
    do
      if (vB)i=0 then vDi:i+31 ← 0x7FFF_FFFF
      else vDi:i+31 ← 0x8000_0000
    VSCRSAT ← 1
  
```


Each single-precision word element in **vB** (word = 32 bits, 4 total elements) is multiplied by  $2^{\text{UIMM}}$ . The product is converted to a signed integer using the rounding mode Round toward Zero. If the intermediate result is greater than  $2^{31}-1$  it saturates to  $2^{31}-1$ ; if it is less than  $-2^{31}$  it saturates to  $-2^{31}$ . A signed-integer result is placed into the corresponding word element of **vD**.

Note that fixed-point integers used by the vector convert instructions can be interpreted as consisting of 32-UIMM integer bits followed by UIMM fraction bits. Also note that the vector convert to fixed-point word instructions support only the rounding mode Round toward Zero. A single-precision number can be converted to a fixed-point integer using any of the other three rounding modes by executing the appropriate vector round to floating-point integer instruction before the vector convert to fixed-point word instruction.

Other registers altered:

- Vector status and control register (VSCR):

Affected: SAT



**Figure 6-43. Convert Floating Point to Integer**

# vctuxs

# vctuxs

Vector Convert to Unsigned Fixed-Point Word Saturate

**vctuxs**                      **vD,vB,UIMM**

04	vD	UIMM	vB	906
0	5 6	10 11	15 16	20 21
				31

```

do i=0 to 127 by 32
  if (vB)i+1:i+8=255 | (vB)i+1:i+8 + UIMM ≤ 254 then
    vDi:i+31 ← CnvtFP32ToUI32Sat((vB)i:i+31 ×fp 2UIMM)
  else
    do
      if (vB)i=0 then vDi:i+31 ← 0xFFFF_FFFF
      else vDi:i+31 ← 0x0000_0000
    VSCRSAT ← 1
  
```

Each single-precision floating-point word element in **vB** (word = 32 bits, 4 total elements) is multiplied by  $2^{UIMM}$ . The product is converted to an unsigned fixed-point integer using the rounding mode Round toward Zero.


If the intermediate result is greater than  $2^{32}-1$  it saturates to  $2^{32}-1$  and if it is less than 0 it saturates to 0.

The unsigned-integer result is placed into the corresponding word element of **vD**.

Other registers altered:

- Vector status and control register (VSCR):

Affected: SAT



**Figure 6-44. Convert Floating Point to Integer**

# vexptefp

Vector 2 Raised to the Exponent Estimate Floating Point

# vexptefp

**vexptefp**

**vD, vB**

04	vD	0 0 0 0 0	vB	458
0	5 6	10 11	15 16	20 21
				31

The single-precision floating-point estimate of 2 raised to the power of each single-precision floating-point element in **vB** is placed into the corresponding element of **vD**.

The estimate has a relative error in precision no greater than one part in 16, that is,

$$\left| \frac{\text{estimate} - 2^x}{2^x} \right| \leq \frac{1}{16}$$


where  $x$  is the value of the element in **vB**. The most significant 12 bits of the estimate's significand are monotonic. Note that the value placed into the element of **vD** may vary between implementations, and between different executions on the same implementation.

Operation with various special values of the element in **vB** is summarized below.

Value	Result
$-\infty$	+0
-0	+1
+0	+1
$+\infty$	$+\infty$
NaN	QNaN

Other registers altered:

- None



**Figure 6-45. Basic One-Source Operands, 32-Bit Elements**

PRELIMINARY

# vlogefp

Vector Log<sub>2</sub> Estimate Floating Point

# vlogefp

**vlogefp**

**vD, vB**

04	vD	0 0 0 0 0	vB	458
0	5 6	10 11	15 16	20 21
				31

The single-precision floating-point estimate of the base 2 logarithm of each single-precision floating-point element in **vB** is placed into the corresponding element of **vD**.

$\log_2$  The estimate has an absolute error in precision (absolute value of the difference between the estimate and the infinitely precise value) no greater than  $2^{-5}$ . The estimate has a relative error in precision no greater than one part in 8, i.e.

$$\left| \frac{\text{estimate} - \log_2(x)}{\log_2(x)} \right| \leq \frac{1}{8}$$

where  $x$  is the value of the element in **vB**, except when  $|x-1| \leq 0.125$ . The most significant 12 bits of the estimate's significand are monotonic. Note that the value placed into the element of **vD** may vary between implementations, and between different executions on the same implementation.


Operation with various special values of the element in **vB** is summarized below.

Value	Result
$-\infty$	QNaN
less than 0	QNaN
$\pm 0$	$-\infty$
$+\infty$	$+\infty$
NaN	QNaN

Other registers altered:

- None





**Figure 6-46. Basic One-Source Operands, 32-Bit Elements**

**PRELIMINARY**

# vmaddfp

Vector Multiply Add Floating Point

# vmaddfp

**vmaddfp**                      **vD,vA,vC,vB**

04	vD	vA	vB	vC	46
0	5 6	10 11	15 16	20 21	26 31

do i=0 to 127 by 32


$\mathbf{vD}_{i:i+31} \leftarrow \text{RndToNearFP32}(((\mathbf{vA})_{i:i+31} \times_{\text{fp}} (\mathbf{vC})_{i:i+31}) +_{\text{fp}} (\mathbf{vB})_{i:i+31})$

Each single-precision floating-point word element in **vA** (word = 32 bits, 4 total elements) is multiplied by the corresponding single-precision floating-point word element in **vC**. The corresponding single-precision floating-point word element in **vB** is added to the product. The result is rounded to the nearest single-precision floating-point number and placed into the corresponding word element of **vD**.

Note that a vector multiply floating-point instruction is not provided. The effect of such an instruction can be obtained by using **vmaddfp** with **vB** containing the value -0 (0x8000\_0000) in each of its four single-precision floating-point word elements. (The value must be -0, not +0, in order to obtain the IEEE-conforming result of -0 when the result of the multiplication is -0.)

Other registers altered:

- None



**Figure 6-47. Multiply Add Floating Point**

**PRELIMINARY**

# vmaxfp

Vector Maximum Floating Point

# vmaxfp

**vmaxfp** **vD,vA,vB**

04	vD	vA	vB	1034
0	5 6	10 11	15 16	20 21
				31

```

do i=0 to 127 by 32
  if (vA)i:i+31 ≥fp (vB)i:i+31
    then vDi:i+31 ← (vA)i:i+31
    else vDi:i+31 ← (vB)i:i+31


```

Each single-precision floating-point word element in **vA** (word = 32 bits, 4 total elements) is compared to the corresponding single-precision floating-point word element in **vB**. The larger of the two single-precision floating-point values is placed into the corresponding word element of **vD**.

The maximum of +0 and -0 is +0. The maximum of any value and a NaN is a QNaN.

Other registers altered:

- None



**Figure 6-48. Basic 2-Source Operands—Four 32-bit Elements**

# vmaxsb

Vector Maximum Signed Byte

# vmaxsb

**vmaxsb**  $\mathbf{vD}, \mathbf{vA}, \mathbf{vB}$

04	$\mathbf{vD}$	$\mathbf{vA}$	$\mathbf{vB}$	258
0	5 6	10 11	15 16	20 21
				31

```

n ← LENGTH(element)
do i=0 to 127 by n
  if ( $\mathbf{vA}$ )i:i+n-1  $\geq_{\text{si}}$  ( $\mathbf{vB}$ )i:i+n-1
    then  $\mathbf{vD}_{i:i+n-1} \leftarrow (\mathbf{vA})_{i:i+n-1}$ 
    else  $\mathbf{vD}_{i:i+n-1} \leftarrow (\mathbf{vB})_{i:i+n-1}$ 


```

For **vmaxsb** each element is a byte.

Each signed-integer element in **vA** is compared to the corresponding signed-integer element in **vB**. The larger of the two signed-integer values is placed into the corresponding element of **vD**.

Other registers altered:

- None



**Figure 6-49. Basic 2-Source Operands , 8-Bit Elements**

# vmaxsh

Vector Maximum Signed Half Word

# vmaxsh

**vmaxsh** **vD,vA,vB**

04	vD	vA	vB	322
0	5 6	10 11	15 16	20 21
				31

```

n ← LENGTH(element)
do i=0 to 127 by n
  if (vA)i:i+n-1 ≥si (vB)i:i+n-1
    then vDi:i+n-1 ← (vA)i:i+n-1
  else vDi:i+n-1 ← (vB)i:i+n-1


```

For **vmaxsh** each element is a half word.

Each signed-integer element in **vA** is compared to the corresponding signed-integer element in **vB**. The larger of the two signed-integer values is placed into the corresponding element of **vD**.

Other registers altered:

- None



**Figure 6-50. Basic 2-Source Operands, Eight 16-bit Elements**

# vmaxsw

Vector Maximum Signed Word

# vmaxsw

**vmaxsw** **vD,vA,vB**

04	vD	vA	vB	386
0	5 6	10 11	15 16	20 21
				31

```

n ← LENGTH(element)
do i=0 to 127 by n
  if (vA)i:i+n-1 ≥si (vB)i:i+n-1
    then vDi:i+n-1 ← (vA)i:i+n-1
  else vDi:i+n-1 ← (vB)i:i+n-1


```

For **vmaxsw** each element is a word.

Each signed-integer element in **vA** is compared to the corresponding signed-integer element in **vB**. The larger of the two signed-integer values is placed into the corresponding element of **vD**.

Other registers altered:

- None




**Figure 6-51. Basic 2-Source Operands—Four 32-bit Elements**

# vmaxub

Vector Maximum Signed Byte

# vmaxub

**vmaxub** **vD,vA,vB**



```

n ← LENGTH(element)
do i=0 to 127 by n
  if (vA)i:i+n-1 ≥ui (vB)i:i+n-1
    then vDi:i+n-1 ← (vA)i:i+n-1
    else vDi:i+n-1 ← (vB)i:i+n-1


```

For **vmaxub** each element is a byte.

Each unsigned-integer element in **vA** is compared to the corresponding unsigned-integer element in **vB**. The larger of the two unsigned-integer values is placed into the corresponding element of **vD**.

Other registers altered:

- None



**Figure 6-52. Basic 2-Source Operands , 8-Bit Elements**



# vmaxuh

Vector Maximum Unsigned Half Word

# vmaxuh

**vmaxuh** **vD,vA,vB**

04	vD	vA	vB	66
0	5 6	10 11	15 16	20 21
				31

```

n ← LENGTH(element)
do i=0 to 127 by n
  if (vA)i:i+n-1 ≥ui (vB)i:i+n-1
    then vDi:i+n-1 ← (vA)i:i+n-1
  else vDi:i+n-1 ← (vB)i:i+n-1


```

For **vmaxuh** each element is a half word.

Each unsigned-integer element in **vA** is compared to the corresponding unsigned-integer element in **vB**. The larger of the two unsigned-integer values is placed into the corresponding element of **vD**.

Other registers altered:

- None




**Figure 6-53. Basic 2-Source Operands, Eight 16-bit Elements**

# vmaxuw

Vector Maximum Unsigned Word

# vmaxuw

**vmaxuw** **vD,vA,vB**



```

n ← LENGTH(element)
do i=0 to 127 by n
  if (vA)i:i+n-1 ≥ui (vB)i:i+n-1
    then vDi:i+n-1 ← (vA)i:i+n-1
  else vDi:i+n-1 ← (vB)i:i+n-1


```

For **vmaxuw** each element is a word.

Each unsigned-integer element in **vA** is compared to the corresponding unsigned-integer element in **vB**. The larger of the two unsigned-integer values is placed into the corresponding element of **vD**.

Other registers altered:

- None



**Figure 6-54. Basic 2-Source Operands—Four 32-bit Elements**

# vmhaddshs

# vmhaddshs

Vector Multiply High and Add Signed Half Word Saturate

**vmhaddshs**      **vD,vA,vB,vC**

04					vD					vA					vB					vC					32																																		
0					5					6					10					11					15					16					20					21					25					26					31				

```

do i=0 to 127 by 16
  prod0-31 ← (vA)i:i+15 ×si (vB)i:i+15
  temp0-16 ← prod0:16 +int SignExtend((vC)i:i+15,17)
  vDi:i+15 ← SItoSIsat(temp0-16,16)

```

Each signed-integer half word element in **vA** is multiplied by the corresponding signed-integer half word element in **vB**, producing a 32-bit signed-integer product. Bits 0-16 of the product are added to the corresponding signed-integer half-word element in **vC**.

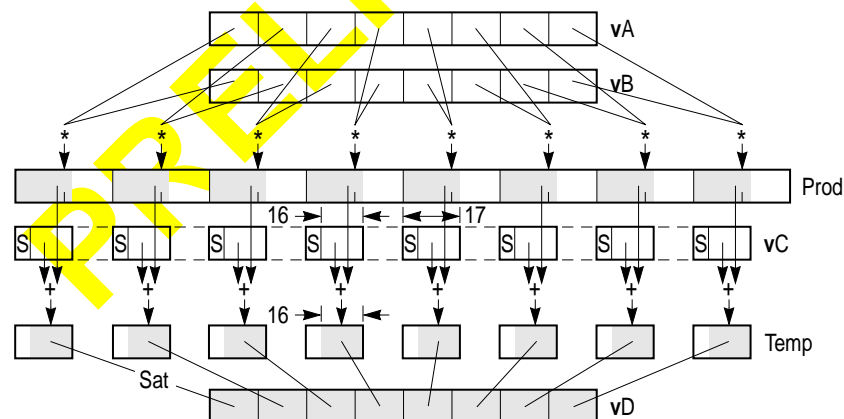
If the intermediate result is greater than  $2^{15}-1$  it saturates to  $2^{15}-1$  and if it is less than  $-2^{15}$  it saturates to  $-2^{15}$ .

The signed-integer result is placed into the corresponding half-word element of **vD**.

Other registers altered:

- Vector status and control register (VSCR):

Affected: SAT



**Figure 6-55. Multiply-High and Add**

# vmhraddshs

# vmhraddshs

Vector Multiply High and Add Signed Half Word Saturate

**vmhraddshs**      **vD,vA,vB,vC**

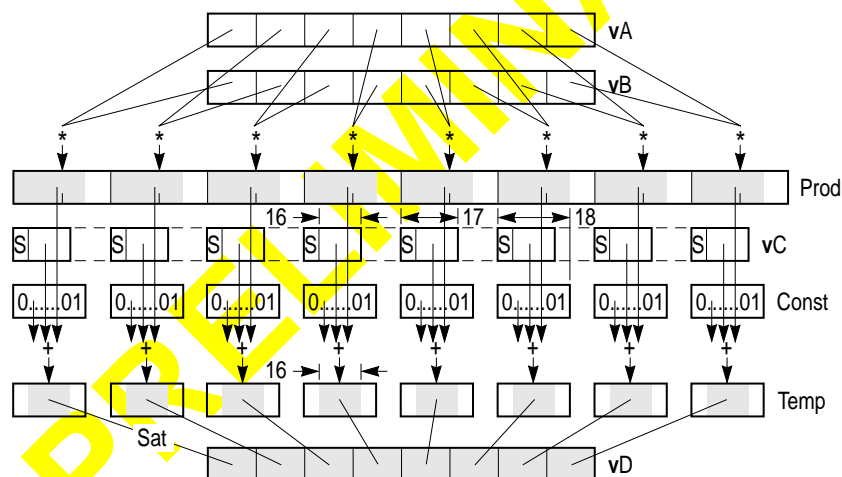
04		vD		vA		vB		vC		33	
0	5	6	10	11	15	16	20	21	25	26	31

```

do i=0,127,16
  prod0-31 ← (vA)i:i+15 *si (vB)i:i+15
  temp0-16 ← (prod0-17 +int (SignExtend((vC)i:i+15,17) || 0b1))0-16
  (vD)i:i+15 ← SItoSIsat(temp0-16,16)
  (VSCRsat) ← (VSCRsat) ∨ ((vD)i:15 saturated)
end

```

The eight 16-bit signed integers in **vA** are multiplied by the eight 16-bit signed integers in **vB**. The intermediate product is rounded to 17 bits of significance and each intermediate rounded product added to the 16-bit signed integers in **vC** after they have been sign extended to 17-bits. The 16-bit saturated result from each of the eight 17-bit sums is placed in **vD**.



**Figure 6-56. Multiply-High Round and Add**

# vminfp

Vector Minimum Floating Point

# vminfp

**vminfp** **vD,vA,vB**


04	vD	vA	vB	1098
0	5 6	10 11	15 16	20 21
				31

```
do i=0,127,32
  (vD)i:i+31 ← FPminimum((vA)i:i+31, (vB)i:i+31)
end
```

The four 32-bit floating-point values in **vA** are compared against the four 32-bit floating-point values in register **rB**. For each of the four pairs of values, the smaller floating-point value within each pair is placed into **vD**.

**vminfp** is sensitive to the sign of 0.0. When both operands are  $\pm 0.0$ :

- $\min(-0.0, \pm 0.0) = \min(\pm 0.0, -0.0) \Rightarrow -0.0$
- $\min(+0.0, +0.0) \Rightarrow +0.0$



**Figure 6-57. Basic 2-Source Operands —32-Bit Elements**

# vminsb

Vector Minimum Signed Byte

# vminsb

**vminsb** **vD,vA,vB**


04	vD	vA	vB	770
0	5 6	10 11	15 16	20 21
				31

```

do i=0,127,N
  if F="u" then (vD)i:i+N-1 ← UIminimum((vA)i:i+N-1, (vB)i:i+N-1)
  if F="s" then (vD)i:i+N-1 ← SIminimum((vA)i:i+N-1, (vB)i:i+N-1)
end

```

The sixteen 8-bit unsigned or signed integers in **vA** are compared against the sixteen 8-bit unsigned or signed integers in register **rB**. For each of the sixteen pairs of values, the larger signed or unsigned integer within each pair is placed into **vD**.



**Figure 6-58. Basic 2-Source Operands , 8-Bit Elements**

# vminsh

Vector Minimum Signed Half Word

# vminsh

**vminsh** **vD,vA,vB**


04	vD	vA	vB	834
0	5 6	10 11	15 16	20 21
				31

```

do i=0,127,N
  if F="u" then (vD)i:i+N-1 ← UIminimum((vA)i:i+N-1, (vB)i:i+N-1)
  if F="s" then (vD)i:i+N-1 ← SIminimum((vA)i:i+N-1, (vB)i:i+N-1)
end

```

The eight 16-bit unsigned or signed integers in **vA** are compared against the eight 16-bit unsigned or signed integers in register **rB**. For each of the eight pairs of values, the larger signed or unsigned integer within each pair is placed into **vD**.



**Figure 6-59. Basic 2-Source Operands , 16-Bit Elements**

# vminsw

Vector Minimum Signed Word

# vminsw

**vminsw** **vD,vA,vB**

04	vD	vA	vB	898
0	5 6	10 11	15 16	20 21
				31

```

do i=0,127,N
  if F="u" then (vD)i:i+N-1 ← UIminimum((vA)i:i+N-1, (vB)i:i+N-1)
  if F="s" then (vD)i:i+N-1 ← SIminimum((vA)i:i+N-1, (vB)i:i+N-1)
end

```

The four 32-bit unsigned or signed integers in **vA** are compared against the four 32-bit unsigned or signed integers in register **vB**. For each of the four pairs of values, the larger signed or unsigned integer within each pair is placed into **vD**.




Figure 6-60. Basic 2-Source Operands , 32-Bit Elements



# vminub

Vector Minimum Unsigned Byte

# vminub

**vminub**                      **vD,vA,vB**


04	vD	vA	vB	514
0	5 6	10 11	15 16	20 21
				31

```

do i=0,127,N
  if F="u" then (vD)i:i+N-1 ← UIminimum((vA)i:i+N-1, (vB)i:i+N-1)
  if F="s" then (vD)i:i+N-1 ← SIminimum((vA)i:i+N-1, (vB)i:i+N-1)
end

```

The sixteen 8-bit unsigned or signed integers in **vA** are compared against the sixteen 8-bit unsigned or signed integers in register **vB**. For each of the sixteen pairs of values, the larger signed or unsigned integer within each pair is placed into **vD**.



**Figure 6-61. Basic 2-Source Operands , 8-bit Elements**

# vminuh

Vector Minimum Unsigned Half Word

# vminuh

**vminuh** **vD,vA,vB**


04	vD	vA	vB	578
0	5 6	10 11	15 16	20 21
				31

```

do i=0,127,N
  if F="u" then (vD)i:i+N-1 ← UIminimum((vA)i:i+N-1, (vB)i:i+N-1)
  if F="s" then (vD)i:i+N-1 ← SIminimum((vA)i:i+N-1, (vB)i:i+N-1)
end

```

The eight 16-bit unsigned or signed integers in **vA** are compared against the eight 16-bit unsigned or signed integers in register **vB**. For each of the eight pairs of values, the larger signed or unsigned integer within each pair is placed into **vD**.



**Figure 6-62. Basic 2-Source Operands , 16-Bit Elements**

# vminuw

Vector Minimum Unsigned Word

# vminuw

**vminuw**                      **vD,vA,vB**


04	vD	vA	vB	642
0	5 6	10 11	15 16	20 21
				31

```

do i=0,127,N
  if F="u" then (vD)i:i+N-1 ← UIminimum((vA)i:i+N-1, (vB)i:i+N-1)
  if F="s" then (vD)i:i+N-1 ← SIminimum((vA)i:i+N-1, (vB)i:i+N-1)
end

```

The four 32-bit unsigned or signed integers in **vA** are compared against the four 32-bit unsigned or signed integers in register **vB**. For each of the four pairs of values, the larger signed or unsigned integer within each pair is placed into **vD**.



**Figure 6-63. Basic 2-Source Operands , 32-Bit Elements**

# vmladduhm

# vmladduhm

Vector Multiply Low and Add Unsigned Half Word Modulo

**vmladduhm**      **vD,vA,vB,vC**

04	vD	vA	vB	vC	34
0	5 6	10 11	15 16	20 21	25 26
					31


```
do i=0 to 127 by 16
  prod0:31 ← (vA)i:i+15 ×ui (vB)i:i+15
  vDi:i+15 ← prod0:31 +int (vC)i:i+15
```

Each unsigned-integer half-word element in **vA** is multiplied by the corresponding unsigned-integer half-word element in **vB**, producing a 32-bit unsigned-integer product. The product is added to the corresponding unsigned-integer half-word element in **vC**. The unsigned-integer result is placed into the corresponding half-word element of **vD**.

Note that **vmladduhm** can be used for unsigned or signed integers.

Other registers altered:

- None



**Figure 6-64. Multiply-Low and Add**

# vmrghb

Vector Merge High Byte

# vmrghb

**vmrghb**  $\mathbf{vD}, \mathbf{vA}, \mathbf{vB}$

04	$\mathbf{vD}$	$\mathbf{vA}$	$\mathbf{vB}$	12
0	5 6	10 11	15 16	20 21
				31

```

n ← LENGTH(element)
do i=0 to 63 by n
     $\mathbf{vD}_{i*2:i*2+n*2-1} \leftarrow (\mathbf{vA})_{i:i+n-1} \parallel (\mathbf{vB})_{i:i+n-1}$ 

```

For **vmrghb** each element is a byte. The elements in the high-order half of  $\mathbf{vA}$  are placed, in the same order, into the even-numbered elements of  $\mathbf{vD}$ . The elements in the high-order half of  $\mathbf{vB}$  are placed, in the same order, into the odd-numbered elements of  $\mathbf{vD}$ .

Other registers altered:

- None




Figure 6-65. Merge High—8-Bit Elements

# vmrghh

Vector Merge High Half word

# vmrghh

**vmrghh**

**vD, vA, vB**

04	vD	vA	vB	76
0	5 6	10 11	15 16	20 21
				31

```


n ← LENGTH(element)
do i=0 to 63 by n
    vDi*2:i*2+n*2-1 ← (vA)i:i+n-1 || (vB)i:i+n-1

```

For **vmrghh** each element is a half word. The elements in the high-order half of **vA** are placed, in the same order, into the even-numbered elements of **vD**. The elements in the high-order half of **vB** are placed, in the same order, into the odd-numbered elements of **vD**.

Other registers altered:

- None



**Figure 6-66. Merge High—16-Bit Elements**

# vmrghw

Vector Merge High Word

# vmrghw

**vmrghw**

**vD, vA, vB**

04	vD	vA	vB	140
0	5 6	10 11	15 16	20 21
				31

```


n ← LENGTH(element)
do i=0 to 63 by n
    vDi*2:i*2+n*2-1 ← (vA)i:i+n-1 || (vB)i:i+n-1

```

For **vmrghw** each element is a word. The elements in the high-order half of **vA** are placed, in the same order, into the even-numbered elements of **vD**. The elements in the high-order half of **vB** are placed, in the same order, into the odd-numbered elements of **vD**.

Other registers altered:

- None



**Figure 6-67. Merge High—32-Bit Elements**

# vmrglb

Vector Merge Low Byte

# vmrglb

**vmrglb**  $\mathbf{vD}, \mathbf{vA}, \mathbf{vB}$

04	$\mathbf{vD}$	$\mathbf{vA}$	$\mathbf{vB}$	268
0	5 6	10 11	15 16	20 21
				31

```


n ← LENGTH(element)
do i=0 to 63 by n
     $\mathbf{vD}_{i*2:i*2+n*2-1} \leftarrow (\mathbf{vA})_{i+64:i+64+n-1} \parallel (\mathbf{vB})_{i+64:i+64+n-1}$ 

```

For **vmrglb** each element is a byte. The elements in the low-order half of  $\mathbf{vA}$  are placed, in the same order, into the even-numbered elements of  $\mathbf{vD}$ . The elements in the low-order half of  $\mathbf{vB}$  are placed, in the same order, into the odd-numbered elements of  $\mathbf{vD}$ .

Other registers altered:

- None



**Figure 6-68. Merge Low—8-Bit Elements**



# vmrglh

Vector Merge Low Half Word

# vmrglh

**vmrglh** **vD,vA,vB**

04	vD	vA	vB	332
0	5 6	10 11	15 16	20 21
				31

```


n ← LENGTH(element)
do i=0 to 63 by n
  vDi*2:i*2+n*2-1 ← (vA)i+64:i+64+n-1 || (vB)i+64:i+64+n-1

```

For **vmrglh** each element is a half word . The elements in the low-order half of **vA** are placed, in the same order, into the even-numbered elements of **vD**. The elements in the low-order half of **vB** are placed, in the same order, into the odd-numbered elements of **vD**.

Other registers altered:

- None



**Figure 6-69. Merge Low—16-Bit Elements**

# vmrglw

Vector Merge Low Word

# vmrglw

**vmrglw**  $\mathbf{vD}, \mathbf{vA}, \mathbf{vB}$

04	$\mathbf{vD}$	$\mathbf{vA}$	$\mathbf{vB}$	396
0	5 6	10 11	15 16	20 21
				31

```


n ← LENGTH(element)
do i=0 to 63 by n
   $\mathbf{vD}_{i*2:i*2+n*2-1} \leftarrow (\mathbf{vA})_{i+64:i+64+n-1} \parallel (\mathbf{vB})_{i+64:i+64+n-1}$ 

```

For **vmrglw** each element is a word. The elements in the low-order half of  $\mathbf{vA}$  are placed, in the same order, into the even-numbered elements of  $\mathbf{vD}$ . The elements in the low-order half of  $\mathbf{vB}$  are placed, in the same order, into the odd-numbered elements of  $\mathbf{vD}$ .

Other registers altered:

- None



**Figure 6-70. Merge Low—32-Bit Elements**

# vmsummbm

Vector Multiply Sum Mixed-Sign Byte Modulo

# vmsummbm

**vmsummbm**      **vD,vA,vB,vC**

04	vD	vA	vB	vC	37
0	5 6	10 11	15 16	20 21	25 26
					31

```

do i=0 to 127 by 32
  temp0-31 ← (vC)i:i+31
  do j=0 to 31 by 8
    prod0-15 ← (vA)i+j:i+j+7 ×sui (vB)i+j:i+j+7
    temp0-31 ← temp0-31 +int SignExtend(prod0-15, 32)
  vDi:i+31 ← temp0-31


```

For each word element in **vC** the following operations are performed in the order shown.

- Each of the four signed-integer byte elements contained in the corresponding word element of **vA** is multiplied by the corresponding unsigned-integer byte element in **vB**, producing a signed-integer product.
- The signed-integer sum of these four products is added to the signed-integer word element in **vC**.
- The signed-integer result is placed into the corresponding word element of **vD**.

Other registers altered:

- None



**Figure 6-71. Multiply-Sum—8-Bit Elements**

# vmsumshm

Vector Multiply Sum Signed Half Word Modulo

# vmsumshm

**vmsumshm**      **vD,vA,vB,vC**

04	vD	vA	vB	vC	40
0	5 6	10 11	15 16	20 21	25 26
					31

```

do i=0 to 127 by 32
  temp0-31 ← (vC)i:i+31
  do j=0 to 31 by 16
    prod0-31 ← (vA)i+j:i+j+15 ×si (vB)i+j:i+j+15
    temp0-31 ← temp0-31 +int prod0-31
  vDi:i+31 ← temp0-31


```

For each word element in **vC** the following operations are performed in the order shown.

- Each of the two signed-integer half-word elements contained in the corresponding word element of **vA** is multiplied by the corresponding signed-integer half-word element in **vB**, producing a signed-integer product.
- The signed-integer sum of these two products is added to the signed-integer word element in **vC**.
- The signed-integer result is placed into the corresponding word element of **vD**.

Other registers altered:

- None



**Figure 6-72. Multiply-Sum—16-Bit Elements**

# vmsumshs

Vector Multiply Sum Signed Half Word Saturate

# vmsumshs

**vmsumshs**      **vD,vA,vB,vC**

04						vD					vA					vB					vC					41				
0		5		6		10		11		15		16		20		21		25		26		31								

```

do i=0 to 127 by 32
  temp0-31 ← SignExtend((vC)i:i+31,34)
  do j=0 to 31 by 16
    prod0-31 ← (vA)i+j:i+j+15 ×si (vB)i+j:i+j+15
    temp0-31 ← temp0-31 +int SignExtend(prod0-31,34)
  vDi:i+31 ← SItoSIsat(temp0-31,32)


```

For each word element in **vC** the following operations are performed in the order shown.

- Each of the two signed-integer half-word elements in the corresponding word element of **vA** is multiplied by the corresponding signed-integer half-word element in **vB**, producing a signed-integer product.
- The signed-integer sum of these two products is added to the signed-integer word element in **vC**.
- If the intermediate result is greater than  $2^{31}-1$  it saturates to  $2^{31}-1$  and if it is less than  $-2^{31}$  it saturates to  $-2^{31}$ .
- The signed-integer result is placed into the corresponding word element of **vD**.

Other registers altered:

- SAT



**Figure 6-73. Multiply-Sum—16-Bit Elements**

# vmsumubm

Vector Multiply Sum Unsigned Byte Modulo

# vmsumubm

**vmsumubm**      **vD,vA,vB,vC**

04	vD	vA	vB	vC	36
0	5 6	10 11	15 16	20 21	25 26
					31

```

do i=0 to 127 by 32
  temp0-31 ← (vC)i:i+31
  do j=0 to 31 by 8
    prod0-15 ← (vA)i+j:i+j+7 ×ui (vB)i+j:i+j+7
    temp0-32 ← temp0-32 +int ZeroExtend(prod0-15, 32)
  vDi:i+31 ← temp0-31


```

For each word element in **vC** the following operations are performed in the order shown.

- Each of the four unsigned-integer byte elements contained in the corresponding word element of **vA** is multiplied by the corresponding unsigned-integer byte element in **vB**, producing an unsigned-integer product.
- The unsigned-integer sum of these four products is added to the unsigned-integer word element in **vC**.
- The unsigned-integer result is placed into the corresponding word element of **vD**.

Other registers altered:

- None



**Figure 6-74. Multiply-Sum—8-Bit Elements**

**PRELIMINARY**

# vmsumuhm

Vector Multiply Sum Unsigned Half Word Modulo

# vmsumuhm

**vmsumuhm**      **vD,vA,vB,vC**

04	vD	A	vB	vC	38
0	5 6	10 11	15 16	20 21	25 26
					31

```

do i=0 to 127 by 32
  temp0-31 ← (vC)i:i+31
  do j=0 to 31 by 16
    prod0-31 ← (vA)i+j:i+j+15 ×ui (vB)i+j:i+j+15
    temp0-31 ← temp0-31 +int prod0-31
  vDi:i+31 ← temp2-33


```

For each word element in **vC** the following operations are performed in the order shown.

- Each of the two unsigned-integer half-word elements contained in the corresponding word element of **vA** is multiplied by the corresponding unsigned-integer half-word element in **vB**, producing an unsigned-integer product.
- The unsigned-integer sum of these two products is added to the unsigned-integer word element in **vC**.
- The unsigned-integer result is placed into the corresponding word element of **vD**.

Other registers altered:

- None



**Figure 6-75. Multiply-Sum, 16-Bit Elements**



# vmsumuhs

# vmsumuhs

Vector Multiply Sum Unsigned Half Word Saturate

**vmsumuhs**      **vD,vA,vB,vC**

04	vD	vA	vB	vC	39
0	5 6	10 11	15 16	20 21	25 26
					31

```

do i=0 to 127 by 32
  temp0-33 ← ZeroExtend((vC)i:i+31,34)
  do j=0 to 31 by 16
    prod0-31 ← (vA)i+j:i+j+15 ×ui (vB)i+j:i+j+15
    temp0-33 ← temp0-33 +int ZeroExtend(prod0-31,34)
  vDi:i+31 ← UItoUISat(temp0-33,32)


```

For each word element in **vC** the following operations are performed in the order shown.

- Each of the two unsigned-integer half-word elements contained in the corresponding word element of **vA** is multiplied by the corresponding unsigned-integer half-word element in **vB**, producing an unsigned-integer product.
- The unsigned-integer sum of these two products is added to the unsigned-integer word element in **vC**.
- The unsigned-integer result is placed into the corresponding word element of **vD**.

Other registers altered:

- SAT



**Figure 6-76. Multiply-Sum—16-Bit Elements**

# vmulesb

# vmulesb

Vector Multiply Even Signed Byte

**vmulesb** **vD,vA,vB**

04	vD	vA	vB	776
0	5 6	10 11	15 16	20 21
				31

```


n ← LENGTH(element)
do i=0 to 127 by n*2
    prod0:n*2-1 ← (vA)i:i+n-1 ×si (vB)i:i+n-1
    vDi:i+n*2-1 ← prod0:n*2-1

```

Each even-numbered signed-integer byte element in **vA** is multiplied by the corresponding signed-integer byte element in **vB**. The eight 16-bit signed-integer products are placed, in the same order, into the eight half-words of **vD**.

Other registers altered:

- None



**Figure 6-77. Multiply, Even Elements, Full-Product, 8-Bit Elements**

# vmulesh

Vector Multiply Even Signed Half Word

# vmulesh

**vmulesh** **vD,vA,vB**

04	vD	vA	vB	840
0	5 6	10 11	15 16	20 21
				31

```


n ← LENGTH(element)
do i=0 to 127 by n*2
  prod0:n*2-1 ← (vA)i:i+n-1 ×si (vB)i:i+n-1
  vDi:i+n*2-1 ← prod0:n*2-1

```

Each even-numbered signed-integer half-word element in **vA** is multiplied by the corresponding signed-integer half-word element in **vB**. The four 32-bit signed-integer products are placed, in the same order, into the four words of **vD**.

Other registers altered:

- None



**Figure 6-78. Multiply, Even Elements, Full-Product, 16-Bit Elements**

# vmuleub

Vector Multiply Even Unsigned Byte

# vmuleub

**vmuleub**  $\mathbf{vD}, \mathbf{vA}, \mathbf{vB}$

04	$\mathbf{vD}$	$\mathbf{vA}$	$\mathbf{vB}$	520
0	5 6	10 11	15 16	20 21
				31

```


do i=0,127,N*2
  prod0:N*2-1 ← ( $\mathbf{vA}$ )i:i+N-1 *ui ( $\mathbf{vB}$ )i:i+N-1
  ( $\mathbf{vD}$ )i:i+N*2-1 ← prod0:N*2-1
end

```

The eight even-numbered 8-bit unsigned or signed integers in  $\mathbf{vA}$  are multiplied by the eight even-numbered 8-bit unsigned or signed integers in  $\mathbf{vB}$ . The eight 16-bit products are placed in  $\mathbf{vD}$ .

Other registers altered:

- None



**Figure 6-79. Multiply, Even Elements, Full-Product, 8-Bit Elements**

# vmuleuh

Vector Multiply Even Unsigned Half Word

# vmuleuh

**vmuleuh** **vD,vA,vB**

04	vD	vA	vB	584
0	5 6	10 11	15 16	20 21
				31

```


do i=0,127,N*2
  prod0:N*2-1 ← (vA)i:i+N-1 *ui (vB)i:i+N-1
  (vD)i:i+N*2-1 ← prod0:N*2-1
end

```

The four even-numbered 16-bit unsigned or signed integers in **vA** are multiplied by the four even-numbered 16-bit unsigned or signed integers in **vB**. The four 32-bit products are placed in **vD**.

Other registers altered:

- None



**Figure 6-80. Multiply, Even Elements, Full-Product, 16-Bit Elements**

# vmulosb

Vector Multiply Odd Signed Byte

# vmulosb

**vmulosb** **vD,vA,vB**

04	vD	vA	vB	264
0	5 6	10 11	15 16	20 21
				31

```

n ← LENGTH(element)
do i=0 to 127 by n*2
    prod0:n*2-1 ← (vA)i+n:i+n+n-1 ×si (vB)i+n:i+n+n-1
    vDi:i+n*2-1 ← prod0:n*2-1

```

Each odd-numbered signed-integer byte element in **vA** is multiplied by the corresponding signed-integer byte element in **vB**. The eight 16-bit signed-integer products are placed, in the same order, into the eight half-words of **vD**.

Other registers altered:

- None




Figure 6-81. Multiply, Odd Elements, Full-Product, 8-Bit Elements

# vmulosh

Vector Multiply Odd Signed Half Word

# vmulosh

**vmulosh** **vD,vA,vB**

04	vD	vA	vB	328
0	5 6	10 11	15 16	20 21
				31

```


n ← LENGTH(element)
do i=0 to 127 by n*2
  prod0:n*2-1 ← (vA)i+n:i+n+n-1 ×si (vB)i+n:i+n+n-1
  vDi:i+n*2-1 ← prod0:n*2-1

```

Each odd-numbered signed-integer half-word element in **vA** is multiplied by the corresponding signed-integer half-word element in **vB**. The four 32-bit signed-integer products are placed, in the same order, into the four words of **vD**.

Other registers altered:

- None



**Figure 6-82. Multiply, Odd Elements, Full-Product, 16-Bit Elements**

# vmuloub

Vector Multiply Odd Unsigned Byte

# vmuloub

**vmuloub** **vD,vA,vB**

04	vD	vA	vB	8
0	5 6	10 11	15 16	20 21
				31

```


n ← LENGTH(element)
do i=0 to 127 by n*2
    prod0:n*2-1 ← (vA)i+n:i+n+n-1 ×ui (vB)i+n:i+n+n-1
    vDi:i+n*2-1 ← prod0:n*2-1

```

Each odd-numbered unsigned-integer byte element in **vA** is multiplied by the corresponding unsigned-integer byte element in **vB**. The eight 16-bit unsigned-integer products are placed, in the same order, into the eight half-words of **vD**.

Other registers altered:

- None



**Figure 6-83. Multiply, Odd Elements, Full-Product, 8-Bit Elements**



# vmulouh

Vector Multiply Odd Unsigned Half Word

# vmulouh

**vmulouh** **vD,vA,vB**

04	vD	vA	vB	72
0	5 6	10 11	15 16	20 21
				31

```


n ← LENGTH(element)
do i=0 to 127 by n*2
  prod0:n*2-1 ← (vA)i+n:i+n+n-1 ×ui (vB)i+n:i+n+n-1
  vDi:i+n*2-1 ← prod0:n*2-1

```

Each odd-numbered unsigned-integer half-word element in **vA** is multiplied by the corresponding unsigned-integer half-word element in **vB**. The four 32-bit unsigned-integer products are placed, in the same order, into the four words of **vD**.

Other registers altered:

- None



**Figure 6-84. Multiply, Odd Elements, Full-Product, 16-Bit Elements**

# vnmsubfp

# vnmsubfp

Vector Negative Multiply-Subtract Floating Point

**vnmsubfp**                      **vD,vA,vC,vB**

04	vD	vA	vB	vC	47
0	5 6	10 11	15 16	20 21	25 26
					31

do i=0 to 127 by 32


$\mathbf{vD}_{i:i+31} \leftarrow -\text{RndToNearFP32}(((\mathbf{vA})_{i:i+31} \times_{\text{fp}} (\mathbf{vC})_{i:i+31}) -_{\text{fp}} (\mathbf{vB})_{i:i+31})$

Each single-precision floating-point word element in **vA** is multiplied by the corresponding single-precision floating-point word element in **vC**. The corresponding single-precision floating-point word element in **vB** is subtracted from the product. The sign of the difference is inverted. The result is rounded to the nearest single-precision floating-point number and placed into the corresponding word element of **vD**.

Other registers altered:

- None

Programming note: The operations negate, absolute value, and negative absolute value can be performed on each single-precision floating-point word element in **vA** by using **vxor**, **vandc**, and **vor** respectively, with **vB** containing the value -0 (0x8000\_0000) in each of its four single-precision floating-point word elements.



# vnor

# vnor

Vector Logical NOR

**vnor** **vD,vA,vB**

04	vD	vA	vB	1284
0	5 6	10 11	15 16	20 21 31

$$vD \leftarrow \neg((vA) \mid (vB))$$

The contents of **vA** are ORed with the contents of **vB** and the complemented result is placed into **vD**.

Other registers altered:

- None

Simplified mnemonics:

**vmr**Vx,Vy equivalent to **vor** Vx,Vy,Vy




Figure 6-85. Basic 2-Source Operands —8-Bit Elements

# vor

Vector Logical OR

# vor

**vor**

**vD,vA,vB**

04	vD	vA	vB	1156
0	5 6	10 11	15 16	20 21 31

$$vD \leftarrow (vA) \mid (vB)$$

The contents of **vA** are ORed with the contents of **vB** and the result is placed into **vD**.

Other registers altered:

- None

Simplified mnemonics

**vmr Vx,Vy** equivalent to **vor Vx,Vy,Vy**




Figure 6-86. Basic 2-Source Operands —8-Bit Elements

# vperm

Vector Permute

# vperm

**vperm** **vD,vA,vB,vC**

04	vD	vA	vB	vC	43
0	5 6	10 11	15 16	20 21	25 26
					31


```
temp0-255 ← (vA) || (vB)
do i=0 to 127 by 8
  b ← (vC)i+3:i+7 || 0b000
  vDi:i+7 ← tempb:b+7
```

Let the source vector be the concatenation of the contents of **vA** followed by the contents of **vB**. For each integer *i* in the range 0–15, the contents of the byte element in the source vector specified in bits 3–7 of byte element *i* in **vC** are placed into byte element *i* of **vD**.

Other registers altered:

- None

Programming note: See the programming notes with the Load Vector for Shift Left and Load Vector for Shift Right instructions for examples of uses of **vperm**.



**Figure 6-87. Permute-8-Bit Elements**

vpkpx

vD,vA,vB

04	vD	vA	vB	782
0	5 6	10 11	15 16	20 21 25 26 31

```

do i=0 to 63 by 16
  vDi ← (vA)i*2+7
  vDi+1:i+5 ← (vA)i*2+8:i*2+12
  vDi+6:i+10 ← (vA)i*2+16:i*2+20
  vDi+11:i+15 ← (vA)i*2+24:i*2+28
  vDi+64 ← (vB)i*2+7
  vDi+65:i+69 ← (vB)i*2+8:i*2+12
  vDi+70:i+74 ← (vB)i*2+16:i*2+20
  vDi+75:i+79 ← (vB)i*2+24:i*2+28

```

The source vector is the concatenation of the contents of **vA** followed by the contents of **vB**. Each word element in the source vector is packed to produce a 16-bit value as described below and placed into the corresponding half-word element of **vD**. A word is packed to 16 bits by concatenating, in order, the following bits.

- bit 7 of the first byte (bit 7 of the word)
- bits 0–4 of the second byte (bits 8–12 of the word)
- bits 0–4 of the third byte (bits 16–20 of the word)
- bits 0–4 of the fourth byte (bits 24–28 of the word)

Other registers altered:

- None

Programming note: Each source word can be considered to be a 32-bit pixel consisting of four 8-bit channels. Each target half-word can be considered to be a 16-bit pixel consisting of one 1-bit channel and three 5-bit channels. A channel can be used to specify the intensity of a particular color, such as red, green, or blue, or to provide other information needed by the application.




Figure 6-88. Pack, 32-Bit Pixels to 1/5/5/5

# vpkshss

Vector Pack Signed Half Word Signed Saturate

# vpkshss

**vpkshss** **vD,vA,vB**

04	vD	vA	vB	398
0	5 6	10 11	15 16	20 21
				31

```

n ← LENGTH(element)
do i=0 to 63 by n/2
  vDi:i+(n/2)-1 ← SItoSIsat((vA)i*2:i*2+n-1,n/2)
  vDi+64:i+64+(n/2)-1 ← SItoSIsat((vB)i*2:i*2+n-1,n/2)


```

Let the source vector be the concatenation of the contents of **vA** followed by the contents of **vB**.

Each signed integer half-word element in the source vector is converted to an 8-bit signed integer. If the value of the element is greater than  $2^7 - 1$  the result saturates to  $2^7 - 1$  and if the value is less than  $-2^7$  the result saturates to  $-2^7$ . The result is placed into the corresponding byte element of **vD**.

Other registers altered:

- SAT



**Figure 6-89. Pack, 16-Bit Elements**

# vpkshus

Vector Pack Signed Half Word Unsigned Saturate

# vpkshus

**vpkshus** **vD,vA,vB**

04	vD	vA	vB	270
0	5 6	10 11	15 16	20 21
				31

```

n ← LENGTH(element)
do i=0 to 63 by n/2
  vDi:i+(n/2)-1 ← SItOUIsat((vA)i*2:i*2+n-1,n/2)
  vDi+64:i+64+(n/2)-1 ← SItOUIsat((vB)i*2:i*2+n-1,n/2)


```

Let the source vector be the concatenation of the contents of **vA** followed by the contents of **vB**.

Each signed integer half-word element in the source vector is converted to an 8-bit unsigned integer. If the value of the element is greater than  $2^8 - 1$  the result saturates to  $2^8 - 1$  and if the value is less than 0 the result saturates to 0. The result is placed into the corresponding byte element of **vD**.

Other registers altered:

- SAT



**Figure 6-90. Pack—16-Bit Elements**



# vpkswus

Vector Pack Signed Word Unsigned Saturate

# vpkswus

**vpkswus** **vD,vA,vB**

04	vD	vA	vB	462
0	5 6	10 11	15 16	20 21
				31

```

n ← LENGTH(element)
do i=0 to 63 by n/2
  vDi:i+(n/2)-1 ← SItoSIsat((vA)i*2:i*2+n-1,n/2)
  vDi+64:i+64+(n/2)-1 ← SItoSIsat((vB)i*2:i*2+n-1,n/2)


```

Let the source vector be the concatenation of the contents of **vA** followed by the contents of **vB**.

Each signed integer word element in the source vector is converted to a 16-bit signed integer. If the value of the element is greater than  $2^{15} - 1$  the result saturates to  $2^{15} - 1$  and if the value is less than  $-2^{15}$  the result saturates to  $-2^{15}$ . The result is placed into the corresponding half-word element of **vD**.

Other registers altered:

- SAT



**Figure 6-91. Pack—32-Bit Elements**

# vpkswus

Vector Pack Signed Word Unsigned Saturate

# vpkswus

**vpkswus** **vD,vA,vB**

04	vD	vA	vB	334
0	5 6	10 11	15 16	20 21
				31

```

n ← LENGTH(element)
do i=0 to 63 by n/2
  vDi:i+(n/2)-1 ← SItOUIsat((vA)i*2:i*2+n-1,n/2)
  vDi+64:i+64+(n/2)-1 ← SItOUIsat((vB)i*2:i*2+n-1,n/2)


```

Let the source vector be the concatenation of the contents of **vA** followed by the contents of **vB**.

Each signed integer word element in the source vector is converted to a 16-bit unsigned integer. If the value of the element is greater than  $2^{16} - 1$  the result saturates to  $2^{16} - 1$  and if the value is less than 0 the result saturates to 0. The result is placed into the corresponding half-word element of **vD**.

Other registers altered:

- SAT



**Figure 6-92. Pack—32-Bit Elements**

# vpkuhum

Vector Pack Unsigned Half Word Unsigned Modulo

# vpkuhum

**vpkuhum**  $\mathbf{vD}, \mathbf{vA}, \mathbf{vB}$

04	$\mathbf{vD}$	$\mathbf{vA}$	$\mathbf{vB}$	14
0	5 6	10 11	15 16	20 21
				31

```

n ← LENGTH(element)
do i=0 to 63 by n/2
   $\mathbf{vD}_{i:i+(n/2)-1} \leftarrow (\mathbf{vA})_{i*2+(n/2):i*2+n-1}$ 
   $\mathbf{vD}_{i+64:i+64+(n/2)-1} \leftarrow (\mathbf{vB})_{i*2+(n/2):i*2+n-1}$ 


```

Let the source vector be the concatenation of the contents of  $\mathbf{vA}$  followed by the contents of  $\mathbf{vB}$ .

The low-order byte of each half-word element in the source vector is placed into the corresponding byte element of  $\mathbf{vD}$ .

Other registers altered:

- None



**Figure 6-93. Pack—16-Bit Elements**

# vpkuhus

Vector Pack Unsigned Half Word Unsigned Saturate

# vpkuhus

**vpkuhus** **vD,vA,vB**

04	vD	vA	vB	142
0	5 6	10 11	15 16	20 21
				31

```

n ← LENGTH(element)
do i=0 to 63 by n/2
  vDi:i+(n/2)-1 ← UItoUISat((vA)i*2:i*2+n-1,n/2)
  vDi+64:i+64+(n/2)-1 ← UItoUISat((vB)i*2:i*2+n-1,n/2)


```

Let the source vector be the concatenation of the contents of **vA** followed by the contents of **vB**.

Each unsigned integer half-word element in the source vector is converted to an 8-bit unsigned integer. If the value of the element is greater than  $2^8 - 1$  the result saturates to  $2^8 - 1$ . The result is placed into the corresponding byte element of **vD**.

Other registers altered:

- SAT



**Figure 6-94. Pack—16-Bit Elements**

# vpkuwum

Vector Pack Unsigned Word Unsigned Modulo

# vpkuwum

**vpkuwum** **vD,vA,vB**

04	vD	vA	vB	78
0	5 6	10 11	15 16	20 21
				31

```

n ← LENGTH(element)
do i=0 to 63 by n/2
  vDi:i+(n/2)-1 ← (vA)i*2+(n/2):i*2+n-1
  vDi+64:i+64+(n/2)-1 ← (vB)i*2+(n/2):i*2+n-1


```

Let the source vector be the concatenation of the contents of **vA** followed by the contents of **vB**.

The low-order half-word of each word element in the source vector is placed into the corresponding half-word element of **vD**.

Other registers altered:

- None



**Figure 6-95. Pack—32-Bit Elements**

# vpkuwus

Vector Pack Unsigned Word Unsigned Saturate

# vpkuwus

**vpkuwus** **vD,vA,vB**

04	vD	vA	vB	206
0	5 6	10 11	15 16	20 21
				31

```

n ← LENGTH(element)
do i=0 to 63 by n/2
  vDi:i+(n/2)-1 ← UItoUISat((vA)i*2:i*2+n-1,n/2)
  vDi+64:i+64+(n/2)-1 ← UItoUISat((vB)i*2:i*2+n-1,n/2)


```

Let the source vector be the concatenation of the contents of **vA** followed by the contents of **vB**.

Each unsigned integer word element in the source vector is converted to a 16-bit unsigned integer. If the value of the element is greater than  $2^{16} - 1$  the result saturates to  $2^{16} - 1$ . The result is placed into the corresponding half-word element of **vD**.

Other registers altered:

- SAT



**Figure 6-96. Pack—32-Bit Elements**

# vrefp

Vector Reciprocal Estimate Floating Point

# vrefp

**vrefp**

**vD, vB**

04	vD	0 0 0 0 0	vB	266
0	5 6	10 11	15 16	20 21
				31

The single-precision floating-point estimate of the reciprocal of each single-precision floating-point element in **vB** is placed into the corresponding element of **vD**.

For results that are not a zero, QNaN, or infinity, the estimate has a relative error in precision no greater than one part in 4096, i.e.

$$\left| \frac{\text{estimate} - 1/x}{1/x} \right| \leq \frac{1}{4096}$$


where  $x$  is the value of the element in **vB**. Note that the value placed into the element of **vD** may vary between implementations, and between different executions on the same implementation.

Operation with various special values of the element in **vB** is summarized below.

Value	Result
$-\infty$	$-0$
$-0$	$-\infty$
$+0$	$+\infty$
$+\infty$	$+0$
NaN	QNaN


Other registers altered:

- None



**Figure 6-97. Basic One-Source Operands—32-Bit Elements**

Vector Round to Floating-Point Integer toward Minus Infinity

**vrfim****vD, vB**

do i=0 to 127 by 32

 $\mathbf{vD}_{i:i+31} \leftarrow \text{RndToFPInt32Floor}((\mathbf{vB})_{i:i+31})$ 

Each single-precision floating-point word element in **vB** is rounded to a single-precision floating-point integer using the rounding mode Round toward -Infinity, and placed into the corresponding word element of **vD**.

Other registers altered:

- None




Figure 6-98. Basic One-Source Operands—32-Bit Elements




# vrfn

# vrfn

Vector Round to Floating-Point Integer Nearest

**vrfn**

**vD, vB**



```
do i=0 to 127 by 32
  vDi:i+31 ← RndToFPInt32Near((vB)i:i+31)
```

Each single-precision floating-point word element in **vB** is rounded to a single-precision floating-point integer using the rounding mode Round to Nearest, and placed into the corresponding word element of **vD**.

Other registers altered:

- None




Figure 6-99. Basic One-Source Operands—32-Bit Elements


# vrrip

# vrrip

Vector Round to Floating-Point Integer toward Positive Infinity

**vrrip**

**vD,vB**



```
do i=0 to 127 by 32
  vDi:i+31 ← RndToFPInt32Ceil((vB)i:i+31)
```

Each single-precision floating-point word element in **vB** is rounded to a single-precision floating-point integer using the rounding mode Round toward +Infinity, and placed into the corresponding word element of **vD**.

Other registers altered:

- None



Figure 6-100. Basic One-Source Operands—32-Bit Elements


# vrfiz

# vrfiz

Vector Round to Floating-Point Integer toward Zero

**vrfiz**

**vD, vB**



```
do i=0 to 127 by 32
  vDi:i+31 ← RndToFPInt32Trunc((vB)i:i+31)
```

Each single-precision floating-point word element in **vB** is rounded to a single-precision floating-point integer using the rounding mode Round toward Zero, and placed into the corresponding word element of **vD**.

Other registers altered:

- None



Figure 6-101. Basic One-Source Operands—32-Bit Elements

# vrlb

Vector Rotate Left Integer Byte

# vrlb

**vrlb**  $\mathbf{vD, vA, vB}$

04	$\mathbf{vD}$	$\mathbf{vA}$	$\mathbf{vB}$	4
0	5 6	10 11	15 16	20 21
				31

```


n ← LENGTH(element)
b ← log2(n)
do i=0 to 127 by n
    sh ← ( $\mathbf{vB}$ )i+(n-b):i+n-1
     $\mathbf{vD}_{i:i+n-1}$  ← ROTL(( $\mathbf{vA}$ )i:i+n-1, sh)

```

Each element is a byte. Let  $n$  be the length of the element. Each element in  $\mathbf{vA}$  is rotated left by the number of bits specified in the low-order  $\log_2(n)$  bits of the corresponding element in  $\mathbf{vB}$ . The result is placed into the corresponding element of  $\mathbf{vD}$ .

Other registers altered:

- None



**Figure 6-102. Rotates and Shifts—8-Bit Elements**

# vrlh

# vrlh

Vector Rotate Left Integer Half Word

**vrlh** **vD,vA,vB**

04	vD	vA	vB	68
0	5 6	10 11	15 16	20 21
				31

```

n ← LENGTH(element)
b ← log2(n)
do i=0 to 127 by n
    sh ← (vB)i+(n-b):i+n-1
    vDi:i+n-1 ← ROTL((vA)i:i+n-1, sh)


```

Each element is a half word .

Let  $n$  be the length of the element. Each element in **vA** is rotated left by the number of bits specified in the low-order  $\log_2(n)$  bits of the corresponding element in **vB**. The result is placed into the corresponding element of **vD**.

Other registers altered:

- None



**Figure 6-103. Rotates and Shifts—16-Bit Elements**

# vrlw

Vector Rotate Left Integer Word

# vrlw

**vrlw** **vD,vA,vB**

04	vD	vA	vB	132
0	5 6	10 11	15 16	20 21
				31

```


n ← LENGTH(element)
b ← log2(n)
do i=0 to 127 by n
    sh ← (vB)i+(n-b):i+n-1
    vDi:i+n-1 ← ROTL((vA)i:i+n-1, sh)

```

Each element is a word. Let  $n$  be the length of the element. Each element in **vA** is rotated left by the number of bits specified in the low-order  $\log_2(n)$  bits of the corresponding element in **vB**. The result is placed into the corresponding element of **vD**.

Other registers altered:

- None



**Figure 6-104. Rotates and Shifts—32-Bit Elements**


# vrsqrtefp

Vector Reciprocal Square Root Estimate Floating Point

# vrsqrtefp

**vrsqrtefp**

**vD, vB**



The single-precision estimate of the reciprocal of the square root of each single-precision element in **vB** is placed into the corresponding word element of **vD**. The estimate has a relative error in precision no greater than one part in 4096, i.e.


$$\left| \frac{\text{estimate} - 1/\sqrt{x}}{1/\sqrt{x}} \right| \leq \frac{1}{4096}$$

where  $x$  is the value of the element in **vB**. Note that the value placed into the element of **vD** may vary between implementations and between different executions on the same implementation. Operation with various special values of the element in **vB** is summarized below.

Value	Result
$-\infty$	QNaN
less than 0	QNaN
$-0$	$-\infty$
$+0$	$+\infty$
$+\infty$	$+0$
NaN	QNaN

Other registers altered:

- None



**Figure 6-105. Basic One-Source Operands, 32-Bit Elements**

# vsel

Vector Conditional Select

# vsel

**vsel**                      **vD,vA,vB,vC**


04	vD	vA	vB	vC	42
0	5 6	10 11	15 16	20 21	25 26
					31

```
do i=0 to 127
  if (vC)i=0 then vDi ← (vA)i
  else vDi ← (vB)i
```

For each bit in **vC** that contains the value 0, the corresponding bit in **vA** is placed into the corresponding bit of **vD**. For each bit in **vC** that contains the value 1, the corresponding bit in **vB** is placed into the corresponding bit of **vD**.

Other registers altered:

- None



**Figure 6-106. Select—1-Bit Elements**



# vsl

Vector Shift Left

# vsl

**vsl** **vD,vA,vB**

04	vD	vA	vB	452
0	5 6	10 11	15 16	20 21
				31

```

sh ← (vB)125-127
t ← 1
do i = 0 to 127 by 8
    t ← t & ((vB)i+5:i+7 = sh)
if t = 1 then vD ← (vA) <<ui sh
else vD ← undefined


```

The contents of **vA** are shifted left by the number of bits specified in **vB**[125–127]. Bits shifted out of bit 0 are lost. Zeros are supplied to the vacated bits on the right. The result is placed into **vD**.

The contents of the low-order three bits of all byte elements in **vB** must be identical to **vB**[125–127]; otherwise the value placed into **vD** is undefined.

Other registers altered:

- None



**Figure 6-107. Vector Shift Left (Shows Shift Count = 6)**

# vsib

Vector Shift Left Integer Byte

# vsib

**vsib**

**vD,vA,vB**

04	vD	vA	vB	260
0	5 6	10 11	15 16	20 21
				31

```

n ← LENGTH(element)
b ← log2(n)
do i=0 to 127 by n
    sh ← (vB)i+(n-b):i+n-1
    vDi:i+n-1 ← (vA)i:i+n-1 <<ui sh


```

Each element is a byte.

Let  $n$  be the length of the element. Each element in **vA** is shifted left by the number of bits specified in the low-order  $\log_2(n)$  bits of the corresponding element in **vB**. Bits shifted out of bit 0 of the element are lost. Zeros are supplied to the vacated bits on the right. The result is placed into the corresponding element of **vD**.

Other registers altered:

- None



**Figure 6-108. Rotates and Shifts—8-Bit Elements**

# vsldoi

Vector Shift Left Double by Octet Immediate

# vsldoi

**vsldoi**                      **vD,vA,vB,SHB**


04				vD				vA				vB				0	SH				44			
0		5	6		10	11		15	16		20	21	22		25	26								31

$$\mathbf{vD} \leftarrow ((\mathbf{vA}) \parallel (\mathbf{vB})) \ll_{ui} (\mathbf{SHB} \parallel 0b000)$$

Let the source vector be the concatenation of the contents of **vA** followed by the contents of **vB**. Bytes **SHB:SHB+15** of the source vector are placed into **vD**.

Other registers altered:

- None



**Figure 6-109. Shift Left Double (Shows Shift Count = 4)**

# vslh

Vector Shift Left Integer Half Word

# vslh

**vslh**

**vD, vA, vB**

04	vD	vA	vB	324
0	5 6	10 11	15 16	20 21
				31

```

n ← LENGTH(element)
b ← log2(n)
do i=0 to 127 by n
    sh ← (vB)i+(n-b):i+n-1
    vDi:i+n-1 ← (vA)i:i+n-1 <<ui sh


```

Each element is a half word .

Let  $n$  be the length of the element. Each element in **vA** is shifted left by the number of bits specified in the low-order  $\log_2(n)$  bits of the corresponding element in **vB**. Bits shifted out of bit 0 of the element are lost. Zeros are supplied to the vacated bits on the right. The result is placed into the corresponding element of **vD**.

Other registers altered:

- None



**Figure 6-110. Rotates and Shifts, 16-Bit Elements**

# vslo

# vslo

Vector Shift Left by Octet

**vslo** **vD,vA,vB**


04	vD	vA	vB	1036
0	5 6	10 11	15 16	20 21
				31

```
shb ← (vB)121-124
vD ← (vA) <<ui (shb || 0b000)
```

The contents of **vA** are shifted left by the number of bytes specified in **vB**[121–124]. Bytes shifted out of byte 0 are lost. Zeros are supplied to the vacated bytes on the right. The result is placed into **vD**.

Other registers altered:

- None



**Figure 6-111. Shift Left Element (Shows Shift Count = 4)**

# vslw

Vector Shift Left Integer Word

# vslw

**vslw**

**vD,vA,vB**

04	vD	vA	vB	388
0	5 6	10 11	15 16	20 21
				31

```

n ← LENGTH(element)
b ← log2(n)
do i=0 to 127 by n
    sh ← (vB)i+(n-b):i+n-1
    vDi:i+n-1 ← (vA)i:i+n-1 <<ui sh


```

Each element is a word.

Let  $n$  be the length of the element. Each element in **vA** is shifted left by the number of bits specified in the low-order  $\log_2(n)$  bits of the corresponding element in **vB**. Bits shifted out of bit 0 of the element are lost. Zeros are supplied to the vacated bits on the right. The result is placed into the corresponding element of **vD**.

Other registers altered:

- None




**Figure 6-112. Rotates and Shifts, 32-Bit Elements**

# vspltb

Vector Splat Byte

# vspltb

**vspltb**                      **vD,vB,UIMM**



```
n ← LENGTH(element)
b ← UIMM*n
do i=0 to 127 by n
    vDi:i+n-1 ← (vB)b:b+n-1
```


Each element is a byte.

The contents of element UIMM in vB are replicated into each element of vD.

Other registers altered:

- None

Programming note: The vector splat instructions can be used in preparation for performing arithmetic for which one source vector is to consist of elements that all have the same value (for example, multiplying all elements of a vector register by a constant).



**Figure 6-113. Splat 8-Bit Element from vB[Element 7] (UIMM=7)**

# vsplth

Vector Splat Half Word

# vsplth

**vsplth**                      **vD,vB,UIMM**

04	vD	UIMM	vB	588
0	5 6	10 11	15 16	20 21
				31

```

n ← LENGTH(element)
b ← UIMM*n
do i=0 to 127 by n
    vDi:i+n-1 ← (vB)b:b+n-1

```


Each element is a half word element is a half word.

The contents of element UIMM in vB are replicated into each element of vD.

Other registers altered:

- None

Programming note: The vector splat instructions can be used in preparation for performing arithmetic for which one source vector is to consist of elements that all have the same value (for example, multiplying all elements of a vector register by a constant).



**Figure 6-114. Splat 16-Bit Contents from vB[Element 1] (UIMM = 1)**



# vspltisb

Vector Splat Immediate Signed Byte

# vspltisb

**vspltisb**                      **vD,SIMM**

04	vD	SIMM	0 0 0 0 0	780
0	5 6	10 11	15 16	20 21 31


```
n ← LENGTH(element)
do i=0 to 127 by n
    vDi:i+n-1 ← SignExtend(SIMM,n)
```

Each element is a byte.

The value of the SIMM field, sign-extended to the length of the element, is replicated into each element of vD.

Other registers altered:

- None



**Figure 6-115. Splat 8-Bit Value from SIMM**

# vspltish

Vector Splat Immediate Signed Half Word

# vspltish

**vspltish**                      **vD,SIMM**

04	vD	SIMM	0 0 0 0 0	844
0	5 6	10 11	15 16	20 21 31

```

n ← LENGTH(element)
do i=0 to 127 by n
    vDi:i+n-1 ← SignExtend(SIMM,n)


```

Each element is a half word.

The value of the SIMM field, sign-extended to the length of the element, is replicated into each element of vD.

Other registers altered:

- None



**Figure 6-116. Splat 16-Bit Value from SIMM**

# vspltisw

Vector Splat Immediate Signed Word

# vspltisw

**vspltisw**                      **vD,SIMM**

04	vD	SIMM	0 0 0 0 0	908
0	5 6	10 11	15 16	20 21
				31


```
n ← LENGTH(element)
do i=0 to 127 by n
    vDi:i+n-1 ← SignExtend(SIMM,n)
```

Each element is a word.

The value of the SIMM field, sign-extended to the length of the element, is replicated into each element of vD.

Other registers altered:

- None



**Figure 6-117. Splat 32-bit value from SIMM**

# vspltw

Vector Splat Word

# vspltw

**vspltw**                      **vD,vB,UIMM**

04	vD	UIMM	vB	652
0	5 6	10 11	15 16	20 21
				31

```

n ← LENGTH(element)
b ← UIMM*n
do i=0 to 127 by n
    vDi:i+n-1 ← (vB)b:b+n-1

```


Each element is a word.

The contents of element UIMM in vB are replicated into each element of vD.

Other registers altered:

- None

Programming note: The Vector Splat instructions can be used in preparation for performing arithmetic for which one source vector is to consist of elements that all have the same value (for example, multiplying all elements of a Vector Register by a constant).



**Figure 6-118. Splat 8-Bit Contents from vB[Element 7] (UIMM = 7)**

# vSr

Vector Shift Right

# vSr

**vSr**

**vD,vA,vB**

04	vD	vA	vB	708
0	5 6	10 11	15 16	20 21 31

```
sh ← (vB)125-127
t ← 1
do i = 0 to 127 by 8
    t ← t & ((vB)i+5:i+7 = sh)
if t = 1 then vD ← (vA) >>ui sh
else vD ← undefined
```

Let  $sh = vB[125-127]$ ;  $sh$  is the shift count in bits ( $0 \leq sh \leq 7$ ). The contents of  $vA$  are shifted right by  $sh$  bits. Bits shifted out of bit 127 are lost. Zeros are supplied to the vacated bits on the left. The result is placed into  $vD$ .

The contents of the low-order three bits of all byte elements in register  $vB$  must be identical to  $vB[125-127]$ ; otherwise the value placed into register  $vD$  is undefined.

Other registers altered:

- None


Programming notes:

A pair of **vslo** and **vsl** or **vsro** and **vSr** instructions, specifying the same shift count register, can be used to shift the contents of a vector register left or right by the number of bits (0–127) specified in the shift count register. The following example shifts the contents of  $vX$  left by the number of bits specified in  $vY$  and places the result into  $vZ$ .

```
vslo    VZ,VX,VY
vsl     VZ,VZ,VY
```

A double-register shift by a dynamically specified number of bits (0–127) can be performed in six instructions. The following example shifts  $(vW) \parallel (vX)$  left by the number of bits specified in  $vY$  and places the high-order 128 bits of the result into  $vZ$ .

```
vslo    t1,VW,VY    #shift high-order reg left
vsl     t1,t1,VY
vsububm t3,V0,VY    #adjust shift count ((V0)=0)
vsro    t2,VX,t3    #shift low-order reg right
vSr     t2,t2,t3
vor     VZ,t1,t2    #merge to get final result
```



**Figure 6-119. Vector Shift Right (Shift Count = 5 Shown)**

PRELIMINARY

# vsrab

# vsrab

Vector Shift Right Algebraic Byte

**vsrab** **vD,vA,vB**

04	vD	vA	vB	772
0	5 6	10 11	15 16	20 21
				31

```

n ← LENGTH(element)
b ← log2(n)
do i=0 to 127 by n
    sh ← (vB)i+(n-b):i+n-1
    vDi:i+n-1 ← (vA)i:i+n-1 >>si sh


```

Each element is a byte.

Let  $n$  be the length of the element. Each element in **vA** is shifted right by the number of bits specified in the low-order  $\log_2(n)$  bits of the corresponding element in **vB**. Bits shifted out of bit  $n-1$  of the element are lost. Bit 0 of the element is replicated to fill the vacated bits on the left. The result is placed into the corresponding element of **vD**.

Other registers altered:

- None



**Figure 6-120. Rotates and Shifts—8-Bit Elements**

# vsrah

# vsrah

Vector Shift Right Algebraic Half Word

**vsrah** **vD,vA,vB**

04	vD	vA	vB	836
0	5 6	10 11	15 16	20 21
				31

```

n ← LENGTH(element)
b ← log2(n)
do i=0 to 127 by n
  sh ← (vB)i+(n-b):i+n-1
  vDi:i+n-1 ← (vA)i:i+n-1 >>si sh


```

Each element is a half word.

Let  $n$  be the length of the element. Each element in **vA** is shifted right by the number of bits specified in the low-order  $\log_2(n)$  bits of the corresponding element in **vB**. Bits shifted out of bit  $n-1$  of the element are lost. Bit 0 of the element is replicated to fill the vacated bits on the left. The result is placed into the corresponding element of **vD**.

Other registers altered:

- None



**Figure 6-121. Rotates and Shifts—16-Bit Elements**



# vsraw

Vector Shift Right Algebraic Word

# vsraw

**vsraw** **vD,vA,vB**

04	vD	vA	vB	900
0	5 6	10 11	15 16	20 21
				31

```

n ← LENGTH(element)
b ← log2(n)
do i=0 to 127 by n
    sh ← (vB)i+(n-b):i+n-1
    vDi:i+n-1 ← (vA)i:i+n-1 >>si sh


```

Each element is a word.

Let  $n$  be the length of the element. Each element in **vA** is shifted right by the number of bits specified in the low-order  $\log_2(n)$  bits of the corresponding element in **vB**. Bits shifted out of bit  $n-1$  of the element are lost. Bit 0 of the element is replicated to fill the vacated bits on the left. The result is placed into the corresponding element of **vD**.

Other registers altered:

- None



**Figure 6-122. Rotates and Shifts—32-Bit Elements**

# vsrb

Vector Shift Right Byte

# vsrb

**vsrb**

**vD,vA,vB**

04	vD	vA	vB	516
0	5 6	10 11	15 16	20 21
				31

```

n ← LENGTH(element)
b ← log2(n)
do i=0 to 127 by n
    sh ← (vB)i+(n-b):i+n-1
    vDi:i+n-1 ← (vA)i:i+n-1 >>ui sh


```

Each element is a byte.

Let  $n$  be the length of the element. Each element in **vA** is shifted right by the number of bits specified in the low-order  $\log_2(n)$  bits of the corresponding element in **vB**. Bits shifted out of bit  $n-1$  of the element are lost. Zeros are supplied to the vacated bits on the left. The result is placed into the corresponding element of **vD**.

Other registers altered:

- None



**Figure 6-123. Rotates and Shifts—8-Bit Elements**

# vsrh

# vsrh

Vector Shift Right Half Word

**vsrh**

**vD,vA,vB**

04	vD	vA	vB	580
0	5 6	10 11	15 16	20 21
				31

```

n ← LENGTH(element)
b ← log2(n)
do i=0 to 127 by n
    sh ← (vB)i+(n-b):i+n-1
    vDi:i+n-1 ← (vA)i:i+n-1 >>ui sh


```

Each element is a half word.

Let  $n$  be the length of the element. Each element in **vA** is shifted right by the number of bits specified in the low-order  $\log_2(n)$  bits of the corresponding element in **vB**. Bits shifted out of bit  $n-1$  of the element are lost. Zeros are supplied to the vacated bits on the left. The result is placed into the corresponding element of **vD**.

Other registers altered:

- None



**Figure 6-124. Rotates and Shifts—16-Bit Elements**

# vsro

Vector Shift Right by Octet

# vsro

**vsro**

**vD, vA, vB**


04	vD	vA	vB	1100
0	5 6	10 11	15 16	20 21
				31

```
shb ← (vB)121-124
vD ← (vA) >>ui (shb || 0b000)
```

The contents of **vA** are shifted right by the number of bytes specified in **vB**[121–124]. Bytes shifted out of **vA** are lost. Zeros are supplied to the vacated bytes on the left. The result is placed into **vD**.

Other registers altered:

- None



**Figure 6-125. Shift Right Element (Shift Count = 5 Shown)**

# Vsrw

Vector Shift Right Word

# Vsrw

**vsrw**

**vD,vA,vB**

04	vD	vA	vB	644
0	5 6	10 11	15 16	20 21
				31

```

n ← LENGTH(element)
b ← log2(n)
do i=0 to 127 by n
    sh ← (vB)i+(n-b):i+n-1
    vDi:i+n-1 ← (vA)i:i+n-1 >>ui sh


```

Each element is a word.

Let  $n$  be the length of the element. Each element in **vA** is shifted right by the number of bits specified in the low-order  $\log_2(n)$  bits of the corresponding element in **vB**. Bits shifted out of bit  $n-1$  of the element are lost. Zeros are supplied to the vacated bits on the left. The result is placed into the corresponding element of **vD**.

Other registers altered:

- None



**Figure 6-126. Rotates and Shifts—8-Bit Elements**

# vsubcuw

Vector Subtract Carryout Unsigned Word

# vsubcuw

**vsubcuw** **vD,vA,vB**

04	vD	vA	vB	1408
0	5 6	10 11	15 16	20 21
				31

```


do i=0 to 127 by 32
  aop0-32 ← ZeroExtend((vA)i:i+31,33)
  bop0-32 ← ZeroExtend((vB)i:i+31,33)
  temp0-32 ← aop0:32 +int -bop0-32 +int 1
  vDi:i+31 ← ZeroExtend(temp0,32)

```

Each unsigned-integer word element in **vB** is subtracted from the corresponding unsigned-integer word element in **vA**. The complement of the borrow out of bit 0 of the 32-bit difference is zero-extended to 32 bits and placed into the corresponding word element of **vD**.

Other registers altered:

- None



# vsubfp

Vector Subtract Floating Point

# vsubfp

**vsubfp**                      **vD,vA,vB**

04	vD	vA	vB	74
0	5 6	10 11	15 16	20 21
				31


do i=0 to 127 by 32

$\mathbf{vD}_{i:i+31} \leftarrow \text{RndToNearFP32}((\mathbf{vA})_{i:i+31} -_{\text{fp}} (\mathbf{vB})_{i:i+31})$

Each single-precision floating-point word element in **vB** is subtracted from the corresponding single-precision floating-point word element in **vA**. The result is rounded to the nearest single-precision floating-point number and placed into the corresponding word element of **vD**.

Other registers altered:

- None



**Figure 6-127. Basic 2-Source Operands —32-Bit Elements**

# vsubsb

Vector Subtract Signed Byte Saturate

# vsubsb

**vsubsb** **vD,vA,vB**

04	vD	vA	vB	1792
0	5 6	10 11	15 16	20 21
				31

```

n ← LENGTH(element)
do i=0 to 127 by n
  aop0:n ← SignExtend((vA)i:i+n-1,n+1)
  bop0:n ← SignExtend((vB)i:i+n-1,n+1)
  temp0:n ← aop0:n +int -bop0:n +int 1
  vDi:i+n-1 ← SItoSI sat(temp0:n,n)

```

Each element is a byte.

Each signed-integer element in **vB** is subtracted from the corresponding signed-integer element in **vA**.

If the intermediate result is greater than  $2^{n-1}-1$  it saturates to  $2^{n-1}-1$  and if it is less than  $-2^{n-1}$  it saturates to  $-2^{n-1}$ , where  $n$  is the length of the element.

The signed-integer result is placed into the corresponding element of **vD**.

Other registers altered:

- SAT




Figure 6-128. Basic 2-Source Operands , 8-Bit Elements



# vsubshs

Vector Subtract Signed Half Word Saturate

# vsubshs

**vsubshs** **vD,vA,vB**

04	vD	vA	vB	1856
0	5 6	10 11	15 16	20 21
				31

```

n ← LENGTH(element)
do i=0 to 127 by n
  aop0:n ← SignExtend((vA)i:i+n-1,n+1)
  bop0:n ← SignExtend((vB)i:i+n-1,n+1)
  temp0:n ← aop0:n +int -bop0:n +int 1
  vDi:i+n-1 ← SItoSIsat(temp0:n,n)

```

Each element is a half word.


Each signed-integer element in **vB** is subtracted from the corresponding signed-integer element in **vA**.

If the intermediate result is greater than  $2^{n-1}-1$  it saturates to  $2^{n-1}-1$  and if it is less than  $-2^{n-1}$  it saturates to  $-2^{n-1}$ , where n is the length of the element.

The signed-integer result is placed into the corresponding element of **vD**.

Other registers altered:

- SAT



**Figure 6-129. Basic 2-Source Operands , 16-Bit Elements**

# vsubsws

Vector Subtract Signed Word Saturate

# vsubsws

**vsubsws** **vD,vA,vB**

04	vD	vA	vB	1920
0	5 6	10 11	15 16	20 21
				31

```

n ← LENGTH(element)
do i=0 to 127 by n
  aop0:n ← SignExtend((vA)i:i+n-1,n+1)
  bop0:n ← SignExtend((vB)i:i+n-1,n+1)
  temp0:n ← aop0:n +int -bop0:n +int 1
  vDi:i+n-1 ← SItoSIsat(temp0:n,n)

```

Each element is a word.


Each signed-integer element in **vB** is subtracted from the corresponding signed-integer element in **vA**.

If the intermediate result is greater than  $2^{n-1}-1$  it saturates to  $2^{n-1}-1$  and if it is less than  $-2^{n-1}$  it saturates to  $-2^{n-1}$ , where  $n$  is the length of the element.

The signed-integer result is placed into the corresponding element of **vD**.

Other registers altered:

- SAT



**Figure 6-130. Basic 2-Source Operands , 32-Bit Elements**

# vsububm

Vector Subtract Unsigned Byte Modulo

# vsububm

**vsububm** **vD,vA,vB**

04	vD	vA	vB	1024
0	5 6	10 11	15 16	20 21
				31

```

n ← LENGTH(element)
do i=0 to 127 by n
    vDi:i+n-1 ← (vA)i:i+n-1 +int ¬(vB)i:i+n-1

```


Each element is a byte.

Each unsigned-integer element in **vB** is subtracted from the corresponding unsigned-integer element in **vA**. The unsigned-integer result is placed into the corresponding element of **vD**.

Other registers altered:

- None

Programming note: The **vsububm**, **vsubuhm** and **vsubuwm** instructions can be used for unsigned or signed integers.



**Figure 6-131. Basic 2-Source Operands , 8-Bit Elements**

# vsububs

Vector Subtract Unsigned Byte Saturate

# vsububs

**vsububs** **vD,vA,vB**

04	vD	vA	vB	1536
0	5 6	10 11	15 16	20 21
				31

```

n ← LENGTH(element)
do i=0 to 127 by n
  aop0:n ← ZeroExtend((vA)i:i+n-1,n+1)
  bop0:n ← ZeroExtend((vB)i:i+n-1,n+1)
  temp0:n ← aop0:n +int -bop0:n +int 1
  vDi:i+n-1 ← SItOUIsat(temp0:n,n)

```

Each element is a byte.


Each unsigned-integer element in **vB** is subtracted from the corresponding unsigned-integer element in **vA**.

If the intermediate result is less than 0 it saturates to 0, where n is the length of the element.

The unsigned-integer result is placed into the corresponding element of **vD**.

Other registers altered:

- SAT



**Figure 6-132. Basic 2-Source Operands , 8-Bit Elements**

# vsubuhm

Vector Subtract Signed Half Word Modulo

# vsubuhm

**vsubuhm** **vD,vA,vB**

04	vD	vA	vB	1088
0	5 6	10 11	15 16	20 21
				31

```

n ← LENGTH(element)
do i=0 to 127 by n
    vDi:i+n-1 ← (vA)i:i+n-1 +int ¬(vB)i:i+n-1

```

Each element is a half word.

Each unsigned-integer element in **vB** is subtracted from the corresponding unsigned-integer element in **vA**. The unsigned-integer result is placed into the corresponding element of **vD**.

Other registers altered:

- None

Programming note: The **vsububm**, **vsubuhm** and **vsubuwm** instructions can be used for unsigned or signed integers.




Figure 6-133. Basic 2-Source Operands , 16-Bit Elements

# vsubuhs

Vector Subtract Signed Half Word Saturate

# vsubuhs

**vsubuhs** **vD,vA,vB**

04	vD	vA	vB	1600
0	5 6	10 11	15 16	20 21
				31

```

n ← LENGTH(element)
do i=0 to 127 by n
  aop0:n ← ZeroExtend((vA)i:i+n-1,n+1)
  bop0:n ← ZeroExtend((vB)i:i+n-1,n+1)
  temp0:n ← aop0:n +int -bop0:n +int 1
  vDi:i+n-1 ← SItOUIsat(temp0:n,n)

```

Each element is a half word.


Each unsigned-integer element in **vB** is subtracted from the corresponding unsigned-integer element in **vA**.

If the intermediate result is less than 0 it saturates to 0, where n is the length of the element.

The unsigned-integer result is placed into the corresponding element of **vD**.

Other registers altered:

- SAT



**Figure 6-134. Basic 2-Source Operands , 16-Bit Elements**

# vsubuwm

Vector Subtract Unsigned Word Modulo

# vsubuwm

**vsubuwm** **vD,vA,vB**

04	vD	vA	vB	1152
0	5 6	10 11	15 16	20 21
				31

```

n ← LENGTH(element)
do i=0 to 127 by n
    vDi:i+n-1 ← (vA)i:i+n-1 +int ¬(vB)i:i+n-1

```

Each element is a word.

Each unsigned-integer element in **vB** is subtracted from the corresponding unsigned-integer element in **vA**. The unsigned-integer result is placed into the corresponding element of **vD**.

Other registers altered:

- None

Programming note: The **vsububm**, **vsubuhm** and **vsubuwm** instructions can be used for unsigned or signed integers.



Figure 6-135. Basic 2-Source Operands , 32-Bit Elements

# vsubuws

Vector Subtract Unsigned Word Saturate

# vsubuws

**vsubuws** **vD,vA,vB**

04	vD	vA	vB	1664
0	5 6	10 11	15 16	20 21 31

```

n ← LENGTH(element)
do i=0 to 127 by n
  aop0:n ← ZeroExtend((vA)i:i+n-1,n+1)
  bop0:n ← ZeroExtend((vB)i:i+n-1,n+1)
  temp0:n ← aop0:n +int -bop0:n +int 1
  vDi:i+n-1 ← SItouIsat(temp0:n,n)

```

Each element is a word.

Each unsigned-integer element in **vB** is subtracted from the corresponding unsigned-integer element in **vA**.

If the intermediate result is less than 0 it saturates to 0, where n is the length of the element.

The unsigned-integer result is placed into the corresponding element of **vD**.

Other registers altered:

- SAT



Figure 6-136. Basic 2-Source Operands , 32-Bit Elements



# vsumsws

Vector Sum Across Signed Word Saturate

# vsumsws

**vsumsws** **vD,vA,vB**


04	vD	vA	vB	1928
0	5 6	10 11	15 16	20 21
				31

```
temp0-34 ← SignExtend((vB)96-127,35)
do i=0 to 127 by 32
    temp0-34 ← temp0-34 +int SignExtend((vA)i:i+31,35)
vD ← 960 || SItoSIsat(temp0-34,32)
```

The signed-integer sum of the four signed-integer word elements in **vA** is added to the signed-integer word element in bits of **vB**[96-127]. If the intermediate result is greater than  $2^{31}-1$  it saturates to  $2^{31}-1$  and if it is less than  $-2^{31}$  it saturates to  $-2^{31}$ . The signed-integer result is placed into bits of **vD**[96-127]. Bits of **vD**[ 0-95] are cleared.

Other registers altered:

- SAT



**Figure 6-137. Sum-Across, 32-Bit Elements**

# vsum2sws

# vsum2sws

Vector Sum Across Partial (1/2) Signed Word Saturate

**vsum2sws** **vD,vA,vB**

04	vD	vA	vB	1672
0	5 6	10 11	15 16	20 21
				31

```


do i=0 to 127 by 64
  temp0:33 ← SignExtend((vB)i+32:i+63,34)
  do j=0 to 63 by 32
    temp0:33 ← temp0:33 +int SignExtend((vA)i+j:i+j+31,34)
  vDi:i+63 ← 320 || SItoSIsat(temp0:33,32)

```

The signed-integer sum of the first two signed-integer word elements in register **vA** is added to the signed-integer word element in **vB**[32–63]. If the intermediate result is greater than  $2^{31}-1$  it saturates to  $2^{31}-1$  and if it is less than  $-2^{31}$  it saturates to  $-2^{31}$ . The signed-integer result is placed into **vD**[32–63]. The signed-integer sum of the last two signed-integer word elements in register **vA** is added to the signed-integer word element in **vB**[96–127]. If the intermediate result is greater than  $2^{31}-1$  it saturates to  $2^{31}-1$  and if it is less than  $-2^{31}$  it saturates to  $-2^{31}$ . The signed-integer result is placed into **vD**[96–127]. The register **vD**[0–31,64–95] are cleared to 0.

Other registers altered:

- SAT



**Figure 6-138. Partial (1/2) Sum Across, 32-Bit Elements**

# vsum4sbs

# vsum4sbs

Vector Sum Across Partial (1/4) Signed Byte Saturate

**vsum4sbs** **vD,vA,vB**

04	vD	vA	vB	1800
0	5 6	10 11	15 16	20 21
				31

```

do i=0 to 127 by 32
  temp0:32 ← SignExtend((vB)i:i+31,33)
  do j=0 to 31 by 8
    temp0:32 ← temp0:32 +int SignExtend((vA)i+j:i+j+7,33)
  vDi:i+31 ← SItoSIsat(temp0:32,32)


```

For each word element in **vB** the following operations are performed in the order shown.

- The signed-integer sum of the four signed-integer byte elements contained in the corresponding word element of register **vA** is added to the signed-integer word element in register **vB**.
- If the intermediate result is greater than  $2^{31}-1$  it saturates to  $2^{31}-1$  and if it is less than  $-2^{31}$  it saturates to  $-2^{31}$ .
- The signed-integer result is placed into the corresponding word element of VT.

Other registers altered:

- SAT



**Figure 6-139. Partial (1/4) Sum-Across, 8-Bit Elements**

# vsum4shs

# vsum4shs

Vector Sum Across Partial (1/4) Signed Half Word Saturate

**vsum4shs** **vD,vA,vB**

04	vD	vA	vB	1608
0	5 6	10 11	15 16	20 21
				31

```

do i=0 to 127 by 32
  temp0:32 ← SignExtend((vB)i:i+31,33)
  do j=0 to 31 by 16
    temp0:32 ← temp0:32 +int SignExtend((vA)i+j:i+j+15,33)
  vDi:i+31 ← SItoSIsat(temp0:32,32)


```

For each word element in register **vB** the following operations are performed, in the order shown.

- The signed-integer sum of the two signed-integer halfword elements contained in the corresponding word element of register **vA** is added to the signed-integer word element in **vB**.
- If the intermediate result is greater than  $2^{31}-1$  it saturates to  $2^{31}-1$  and if it is less than  $-2^{31}$  it saturates to  $-2^{31}$ .
- The signed-integer result is placed into the corresponding word element of **vD**.

Other registers altered:

- SAT



**Figure 6-140. Partial (1/4) Sum-Across, 8-Bit Elements**

# vsum4ubs

# vsum4ubs

Vector Sum Across Partial (1/4) Unsigned Byte Saturate

**vsum4ubs** **vD,vA,vB**

04	vD	vA	vB	1544
0	5 6	10 11	15 16	20 21
				31

```

do i=0 to 127 by 32
  temp0:32 ← ZeroExtend((vB)i:i+31,33)
  do j=0 to 31 by 8
    temp0:32 ← temp0:32 +int ZeroExtend((vA)i+j:i+j+7,33)
  vDi:i+31 ← UItoUISat(temp0:32,32)


```

For each word element in **vB** the following operations are performed in the order shown.

- The unsigned-integer sum of the four unsigned-integer byte elements contained in the corresponding word element of register **vA** is added to the unsigned-integer word element in register **vB**.
- If the intermediate result is greater than  $2^{32}-1$  it saturates to  $2^{32}-1$ .
- The unsigned-integer result is placed into the corresponding word element of **vD**.

Other registers altered:

- SAT



**Figure 6-141. Partial (1/4) Sum-Across, 8-Bit Elements**


# vupkhp<sub>x</sub>

Vector Unpack High Pixel16

# vupkhp<sub>x</sub>

**vupkhp<sub>x</sub>**

**vD, vB**



```

do i=0 to 63 by 16
  vDi*2:i*2+7 ← SignExtend((vB)i, 8)
  vDi*2+8:i*2+15 ← ZeroExtend((vB)i+1:i+5, 8)
  vDi*2+16:i*2+23 ← ZeroExtend((vB)i+6:i+10, 8)
  vDi*2+24:i*2+31 ← ZeroExtend((vB)i+11:i+15, 8)

```

Each halfword element in the high-order half of register **vB** is unpacked to produce a 32-bit value as described below and placed, in the same order, into the four words of **vD**.


A halfword is unpacked to 32 bits by concatenating, in order, the results of the following operations.

- sign-extend bit 0 of the halfword to 8 bits
- zero-extend bits 1–5 of the halfword to 8 bits
- zero-extend bits 6–10 of the halfword to 8 bits
- zero-extend bits 11–15 of the halfword to 8 bits

Other registers altered:

- None

The source and target elements can be considered to be 16-bit and 32-bit "pixels" respectively, having the formats described in the programming note for the Vector Pack Pixel instruction.



**Figure 6-142. Unpack High, 16-Bit Pixels**

**PRELIMINARY**


# vupkhsb

Vector Unpack High Signed Byte

# vupkhsb

**vupkhsb**

**vD, vB**



```


n ← LENGTH(element)
do i=0 to 63 by n
    vDi*2:i*2+n*2-1 ← SignExtend((vB)i:i+n-1, n*2)

```

Each signed integer byte element in the high-order half of register **vB** is sign-extended to produce a 16-bit signed integer and placed, in the same order, into the eight halfwords of register **vD**.

Other registers altered:

- None



**Figure 6-143. Unpack High, 8-Bit Elements**




# vupkhsh

Vector Unpack High Signed Half Word

# vupkhsh

**vupkhsh**

**vD,vB**




```
n ← LENGTH(element)
do i=0 to 63 by n
  vDi*2:i*2+n*2-1 ← SignExtend((vB)i:i+n-1,n*2)
```

Each signed integer halfword element in the high-order half of register **vB** is sign-extended to produce a 32-bit signed integer and placed, in the same order, into the four words of register **vD**.

Other registers altered:

- None



**Figure 6-144. Unpack High, 16-Bit Elements**


# vupklpx

Vector Unpack Low Pixel16

# vupklpx

**vupklpx**

**vD,vB**



```

do i=0 to 63 by 16
  vDi*2:i*2+7 ← SignExtend((vB)i+64,8)
  vDi*2+8:i*2+15 ← ZeroExtend((vB)i+65:i+69,8)
  vDi*2+16:i*2+23 ← ZeroExtend((vB)i+70:i+74,8)
  vDi*2+24:i*2+31 ← ZeroExtend((vB)i+75:i+79,8)

```

Each halfword element in the low-order half of register **vB** is unpacked to produce a 32-bit value as described below and placed, in the same order, into the four words of register **vD**.


A halfword is unpacked to 32 bits by concatenating, in order, the results of the following operations.

- sign-extend bit 0 of the halfword to 8 bits
- zero-extend bits 1–5 of the halfword to 8 bits
- zero-extend bits 6–10 of the halfword to 8 bits
- zero-extend bits 11–15 of the halfword to 8 bits

Other registers altered:

- None

Programming note: Notice that the unpacking done by the Vector Unpack Pixel instructions does not reverse the packing done by the Vector Pack Pixel instruction. Specifically, if a 16-bit pixel is unpacked to a 32-bit pixel which is then packed to a 16-bit pixel, the resulting 16-bit pixel will not, in general, be equal to the original 16-bit pixel (because, for each channel except the first, Vector Unpack Pixel inserts high-order bits while Vector Pack Pixel discards low-order bits).



**Figure 6-145. Unpack Low—16-Bit Pixels**


# vupklsb

Vector Unpack Low Signed Byte

# vupklsb

**vupklsb**

**vD,vB**



```


n ← LENGTH(element)
do i=0 to 63 by n
    vDi*2:i*2+n*2-1 ← SignExtend((vB)i+64:i+64+n-1,n*2)

```

Each signed integer byte element in the low-order half of register **vB** is sign-extended to produce a 16-bit signed integer and placed, in the same order, into the eight halfwords of register **vD**.

Other registers altered:

- None



**Figure 6-146. Unpack Low—8-Bit Elements**


# vupklsh

Vector Unpack Low Signed Half Word

# vupklsh

**vupklsh**

**vD, vB**




```

n ← LENGTH(element)
do i=0 to 63 by n
    vDi*2:i*2+n*2-1 ← SignExtend((vB)i+64:i+64+n-1, n*2)
    
```

Each signed integer half word element in the low-order half of register **vB** is sign-extended to produce a 32-bit signed integer and placed, in the same order, into the four words of register **vD**.

Other registers altered:

- None




**Figure 6-147. Unpack Low—16-Bit Elements**

# vxor

Vector Logical XOR

# vxor

**vxor** **vD,vA,vB**



$$vD \leftarrow (vA) \oplus (vB)$$

The contents of **vA** are XORed with the contents of register **vB** and the result is placed into register **vD**.

Other registers altered:

- None




Figure 6-148

**PRELIMINARY**

# Appendix A

## Altivec Instruction Set Listings

This appendix lists the instruction set for the Altivec™ technology. Instructions are sorted by mnemonic, opcode, and form. Also included in this appendix is a quick reference table that contains general information, such as the architecture level, privilege level, and form, and indicates if the instruction is optional.

Note that split fields, which represent the concatenation of sequences from left to right, are shown in lowercase.

### A.1 Instructions Sorted by Mnemonic

Table A-1 lists the instructions implemented in the Altivec architecture in alphabetical order by mnemonic.

Key:

 Reserved bits

**Table A-1. Complete Instruction List Sorted by Mnemonic**

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
<b>dst</b>	31	T	0	0	STRM					A					B													0
<b>dstst</b>	31	T	0	0	STRM					A					B													0
<b>dststt</b>	31	1	0	0	tag					A					B					11						22		0
<b>dstt</b>	31	1	0	0	tag					A					B													0
<b>dss</b>	31	A	0	0	STRM	0	0	0	0	0	0	0	0	0	0	0	0											0
<b>dssall</b>	31	A	0	0	STRM	0	0	0	0	0	0	0	0	0	0	0	0											0
<b>lvebx</b>	31				vD					A					B													0
<b>lvehx</b>	31				vD					A					B													0
<b>lvewx</b>	31				vD					A					B													0
<b>lvsl</b>	31				vD					A					B													0
<b>lvsr</b>	31				vD					A					B													0
<b>lvx</b>	31				vD					A					B													0

Name 0 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 31

lvxl	31	vD	A				B				359				0	
mfvscr	04	vD	0	0	0	0	0	0	0	0	0	1540				0
mtvscr	04	///	0	0	0	0	0	vD				1604				0
stvebx	31	S	A				B				135				0	
stvehx	31	S	A				B				167				0	
stviewx	31	S	A				B				199				0	
stvx	31	S	A				B				231				0	
stvxl	31	S	A				B				487				0	
vaddcuw	04	vD	vA				vB				384				0	
vaddfp	04	vD	vA				vB				10				0	
vaddsbs	04	vD	vA				vB				768				0	
vaddshs	04	vD	vA				vB				832				0	
vaddsws	04	vD	vA				vB				896				0	
vaddubm	04	vD	vA				vB				0				0	
vaddubs	04	vD	vA				vB				512				0	
vadduhm	04	vD	vA				vB				64				0	
vadduhs	04	vD	vA				vB				576				0	
vadduwm	04	vD	vA				vB				128				0	
vadduws	04	vD	vA				vB				640				0	
vand	04	vD	vA				vB				1028				0	
vandc	04	vD	vA				vB				1092				0	
vavgsb	04	vD	vA				vB				1282				0	
vavgsh	04	vD	vA				vB				1346				0	
vavgsw	04	vD	vA				vB				1410				0	
vavgub	04	vD	vA				vB				1026				0	
vavguh	04	vD	vA				vB				1090				0	
vavguw	04	vD	vA				vB				1154				0	
vcfsx	04	vD	UIMM				vB				842					
vcfux	04	vD	UIMM				vB				778				0	
vcmpbfpx	04	vD	vA				vB				Rc	966				
vcmpeqfx	04	vD	vA				vB				Rc	198				
vcmpequbx	04	vD	vA				vB				Rc	6				
vcmpequhx	04	vD	vA				vB				Rc	70				



Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
<b>vcmpequwx</b>	04				vD					vA					vB			Rc										134
<b>vcmpgefp</b>	04				vD					vA					vB			Rc										454
<b>vcmpgtfp</b>	04				vD					vA					vB			Rc										710
<b>vcmpgtsbx</b>	04				vD					vA					vB			Rc										774
<b>vcmpgtshx</b>	04				vD					vA					vB			Rc										838
<b>vcmpgtswx</b>	04				vD					vA					vB			Rc										902
<b>vcmpgtubx</b>	04				vD					vA					vB			Rc										518
<b>vcmpgtuhx</b>	04				vD					vA					vB			Rc										582
<b>vcmpgtuwx</b>	04				vD					vA					vB			Rc										646
<b>vctxs</b>	04				vD					UIMM					vB													970
<b>vctuxs</b>	04				vD					UIMM					vB													906
<b>vexptfp</b>	04				vD				0	0	0	0	0		vB													458
<b>vlogefp</b>	04				vD				0	0	0	0	0		vB													458
<b>vmaddfp</b>	04				vD					vA					vB				vC									46
<b>vmaxfp</b>	04				vD					vA					vB													1034
<b>vmaxsb</b>	04				vD					vA					vB													258
<b>vmaxsh</b>	04				vD					vA					vB													322
<b>vmaxsw</b>	04				vD					vA					vB													386
<b>vmaxub</b>	04				vD					vA					vB													2
<b>vmaxuh</b>	04				vD					vA					vB													66
<b>vmaxuw</b>	04				vD					vA					vB													130
<b>vmhaddshs</b>	04				vD					vA					vB				vC									32
<b>vmhraddshs</b>	04				vD					vA					vB				vC									33
<b>vminfp</b>	04				vD					vA					vB													1098
<b>vminsb</b>	04				vD					vA					vB													770
<b>vminsh</b>	04				vD					vA					vB													834
<b>vminsw</b>	04				vD					vA					vB													898
<b>vminub</b>	04				vD					vA					vB													514
<b>vminuh</b>	04				vD					vA					vB													578
<b>vminuw</b>	04				vD					vA					vB													642
<b>vmladduhm</b>	04				vD					vA					vB				vC									34
<b>vmrghb</b>	04				vD					vA					vB													12
<b>vmrghh</b>	04				vD					vA					vB													76

Name 0 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 31

vmrghw	04	vD	vA	vB	140			
vmrglb	04	vD	vA	vB	268			
vmrglh	04	vD	vA	vB	332			
vmrglw	04	vD	vA	vB	396			
vmsummbm	04	vD	vA	vB	vC	37		
vmsumshm	04	vD	vA	vB	vC	40		
vmsumshs	04	vD	vA	vB	vC	41		
vmsumubm	04	vD	vA	vB	vC	36		
vmsumuhm	04	vD	vA	vB	vC	38		
vmsumuhs	04	vD	vA	vB	vC	39		
vmulesb	04	vD	vA	vB	776			
vmulesh	04	vD	vA	vB	840			
vmuleub	04	vD	vA	vB	520			
vmuleuh	04	vD	vA	vB	584			
vmulosb	04	vD	vA	vB	264			
vmulosh	04	vD	vA	vB	328			
vmuloub	04	vD	vA	vB	8			
vmulouh	04	vD	vA	vB	72			
vnmsubfp	04	vD	vA	vB	vC	47		
vnor	04	vD	vA	vB	1284			
vor	04	vD	vA	vB	1156			
vperm	04	vD	vA	vB	vC	43		
vpkpx	04	vD	vA	vB	12	782		
vpkshss	04	vD	vA	vB	398			
vpkshus	04	vD	vA	vB	270			
vpkswss	04	vD	vA	vB	462			
vpkuhum	04	vD	vA	vB	14			
vpkuhus	04	vD	vA	vB	142			
vpkuwum	04	vD	vA	vB	78			
vpkuwus	04	vD	vA	vB	206			
vrefp	04	vD	0	0	0	0	vB	266
vrfim	04	vD	0	0	0	0	vB	714
vrfin	04	vD	0	0	0	0	vB	522

Name 0 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 31

vrflp	04	vD	0	0	0	0	0	vB	650				
vrflz	04	vD	0	0	0	0	0	vB	586				
vrllb	04	vD	vA					vB	4				
vrllh	04	vD	vA					vB	68				
vrllw	04	vD	vA					vB	132				
vrsqrtefp	04	vD	0	0	0	0	0	vB	330				
	04	vD	vA					vB	vC	42			
vsl	04	vD	vA					vB	452				
vsllb	04	vD	vA					vB	260				
vsldoi	04	vD	vA					vB	0	SH	44		
vsllh	04	vD	vA					vB	324				
vsllw	04	vD	vA					vB	1036				
vsllw	04	vD	vA					vB	388				
vspltb	04	vD	UIMM					vB	524				
vsplth	04	vD	UIMM					vB	588				
vspltisb	04	vD	SIMM					vB	780				
vspltish	04	vD	SIMM					0	0	0	0	0	844
vspltisw	04	vD	SIMM					0	0	0	0	0	908
vspltw	04	vD	UIMM					vB	652				
vsr	04	vD	vA					vB	708				
vsrab	04	vD	vA					vB	772				
vsrah	04	vD	vA					vB	836				
vsraw	04	vD	vA					vB	900				
vsrb	04	vD	vA					vB	516				
vsrh	04	vD	vA					vB	580				
vsro	04	vD	vA					vB	1100				
vsrw	04	vD	vA					vB	644				
vsubcuw	04	vD	vA					vB	1408				
vsubfp	04	vD	vA					vB	74				
vsubsbbs	04	vD	vA					vB	1792				
vsubshs	04	vD	vA					vB	1856				
vsubsws	04	vD	vA					vB	1920				
vsububm	04	vD	vA					vB	1024				



## A.2 Instructions Sorted by Opcode

Table A-2 lists the AltiVec instructions grouped by opcode.

Key:



Reserved bits

**Table A-2. Instructions Sorted by Opcode**

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
<b>vmhaddshs</b>	000100				vD					vA					vB					vC								32
<b>vmhraddshs</b>	000100				vD					vA					vB					vC								33
<b>vmladduhm</b>	000100				vD					vA					vB					vC								34
<b>vmsumubm</b>	000100				vD					vA					vB					vC								36
<b>vmsummbm</b>	000100				vD					vA					vB					vC								37
<b>vmsumuhm</b>	000100				vD					A					vB					vC								38
<b>vmsumuhs</b>	000100				vD					vA					vB					vC								39
<b>vmsumshm</b>	000100				vD					vA					vB					vC								40
<b>vmsumshs</b>	000100				vD					vA					vB					vC								41
<b>vsel</b>	000100				vD					vA					vB					vC								42
<b>vperm</b>	000100				vD					vA					vB					vC								43
<b>vsldoi</b>	000100				vD					vA					vB		0			SH								44
<b>vmaddfp</b>	000100				vD					vA					vB													46
<b>vnmsubfp</b>	000100				vD					vA					vB					vC								47
<b>vaddubm</b>	000100				vD					vA					vB													0
<b>vadduhm</b>	000100				vD					vA					vB													64
<b>vadduwm</b>	000100				vD					vA					vB													128
<b>vaddcuw</b>	000100				vD					vA					vB													384
<b>vaddubs</b>	000100				vD					vA					vB													512
<b>vadduhs</b>	000100				vD					vA					vB													576
<b>vadduws</b>	000100				vD					vA					vB													640
<b>vaddsbs</b>	000100				vD					vA					vB													768
<b>vaddshs</b>	000100				vD					vA					vB													832
<b>vaddsws</b>	000100				vD					vA					vB													896
<b>vsububm</b>	000100				vD					vA					vB													1024
<b>vsubuhm</b>	000100				vD					vA					vB													1088
<b>vsubuwm</b>	000100				vD					vA					vB													1152
<b>vsubcuw</b>	000100				vD					vA					vB													1408
<b>vsububs</b>	000100				vD					vA					vB													1536

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
vsubuhs	000100				vD					vA					vB														1600
vsubuws	000100				vD					vA					vB														1664
vsubsb	000100				vD					vA					vB														1792
vsubshs	000100				vD					vA					vB														1856
vsubsws	000100				vD					vA					vB														1920
vmaxub	000100				vD					vA					vB														2
vmaxuh	000100				vD					vA					vB														66
vmaxuw	000100				vD					vA					vB														130
vmaxsb	000100				vD					vA					vB														258
vmaxsh	000100				vD					vA					vB														322
vmaxsw	000100				vD					vA					vB														386
vminub	000100				vD					vA					vB														514
vminuh	000100				vD					vA					vB														578
vminuw	000100				vD					vA					vB														642
vminsb	000100				vD					vA					vB														770
vminsh	000100				vD					vA					vB														834
vminsw	000100				vD					vA					vB														898
vavgub	000100				vD					vA					vB														1026
vavguh	000100				vD					vA					vB														1090
vavguw	000100				vD					vA					vB														1154
vavgsb	000100				vD					vA					vB														1282
vavgsh	000100				vD					vA					vB														1346
vavgsw	000100				vD					vA					vB														1410
vrlb	000100				vD					vA					vB														4
vrlh	000100				vD					vA					vB														68
vrlw	000100				vD					vA					vB														132
vs1b	000100				vD					vA					vB														260
vs1h	000100				vD					vA					vB														324
vs1w	000100				vD					vA					vB														388
vsl	000100				vD					vA					vB														452
vsrb	000100				vD					vA					vB														516
vsrh	000100				vD					vA					vB														580
vsrw	000100				vD					vA					vB														644

Name     0                    5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 31

<b>vsr</b>	000100	vD	vA	vB	708		
<b>vsrab</b>	000100	vD	vA	vB	772		
<b>vsrah</b>	000100	vD	vA	vB	836		
<b>vsraw</b>	000100	vD	vA	vB	900		
<b>vand</b>	000100	vD	vA	vB	1028		
<b>vandc</b>	000100	vD	vA	vB	1092		
<b>vor</b>	000100	vD	vA	vB	1156		
<b>vxor</b>	000100	vD	vA	vB	1220		
<b>vnor</b>	000100	vD	vA	vB	1284		
<b>mfvscr</b>	000100	vD	0 0 0 0 0	0 0 0 0 0	1540		0
<b>mtvscr</b>	000100	0 0 0 0 0	0 0 0 0 0	vB	1604		0
<b>vcmpequbx</b>	000100	vD	vA	vB	R c	6	
<b>vcmpequhx</b>	000100	vD	vA	vB	R c	70	
<b>vcmpequwx</b>	000100	vD	vA	vB	R c	134	
<b>vcmpeqfpx</b>	000100	vD	vA	vB	R c	198	
<b>vcmpgefpx</b>	000100	vD	vA	vB	R c	454	
<b>vcmpgtubx</b>	000100	vD	vA	vB	R c	518	
<b>vcmpgtuhx</b>	000100	vD	vA	vB	R c	582	
<b>vcmpgtuwx</b>	000100	vD	vA	vB	R c	646	
<b>vcmpgtfpx</b>	000100	vD	vA	vB	R c	710	
<b>vcmpgtsbx</b>	000100	vD	vA	vB	R c	774	
<b>vcmpgtshx</b>	000100	vD	vA	vB	R c	838	
<b>vcmpgtswx</b>	000100	vD	vA	vB	R c	902	
<b>vcmpbfpx</b>	000100	vD	vA	vB	R c	966	
<b>vmuloub</b>	000100	vD	vA	vB	8		
<b>vmulouh</b>	000100	vD	vA	vB	72		
<b>vmulosb</b>	000100	vD	vA	vB	264		

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	
vmulosh	000100				vD					vA					vB														328
vmuleub	000100				vD					vA					vB														520
vmuleuh	000100				vD					vA					vB														584
vmulesb	000100				vD					vA					vB														776
vmulesh	000100				vD					vA					vB														840
vsum4ubs	000100				vD					vA					vB														1544
vsum4sbs	000100				vD					vA					vB														1800
vsum4shs	000100				vD					vA					vB														1608
vsum2sws	000100				vD					vA					vB														1672
vsumsws	000100				vD					vA					vB														1928
vaddfp	000100				vD					vA					vB														10
vsubfp	000100				vD					vA					vB														74
vrefp	000100				vD					0 0 0 0 0					vB														266
vrsqrtefp	000100				vD					0 0 0 0 0					vB														330
vexptefp	000100				vD					0 0 0 0 0					vB														458
vlogefp	000100				vD					0 0 0 0 0					vB														458
vrfin	000100				vD					0 0 0 0 0					vB														522
vrfiz	000100				vD					0 0 0 0 0					vB														586
vrfip	000100				vD					0 0 0 0 0					vB														650
vrfim	000100				vD					0 0 0 0 0					vB														714
vcfux	000100				vD					UIMM					vB														778
vcfsx	000100				vD					UIMM					vB														842
vctuxs	000100				vD					UIMM					vB														906
vctxsx	000100				vD					UIMM					vB														970
vmaxfp	000100				vD					vA					vB														1034
vminfp	000100				vD					vA					vB														1098
vmrghb	000100				vD					vA					vB														12
vmrghh	000100				vD					vA					vB														76
vmrghw	000100				vD					vA					vB														140
vmrglb	000100				vD					vA					vB														268
vmrglh	000100				vD					vA					vB														332
vmrglw	000100				vD					vA					vB														396
vspltb	000100				vD					UIMM					vB														524



Name 0 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 31

vsplth	000100	vD	UIMM	vB	588						
vspltw	000100	vD	UIMM	vB	652						
vspltisb	000100	vD	SIMM	0 0 0 0 0	780						
vspltish	000100	vD	SIMM	0 0 0 0 0	844						
vspltisw	000100	vD	SIMM	0 0 0 0 0	908						
vslo	000100	vD	vA	vB	1036						
vsro	000100	vD	vA	vB	1100						
vpkuhum	000100	vD	vA	vB	14						
vpkuwum	000100	vD	vA	vB	78						
vpkuhus	000100	vD	vA	vB	142						
vpkuwus	000100	vD	vA	vB	206						
vpkshus	000100	vD	vA	vB	270						
vpkswus	000100	vD	vA	vB	334						
vpkshss	000100	vD	vA	vB	398						
vpkswss	000100	vD	vA	vB	462						
vupkhsb	000100	vD	0 0 0 0 0	vB	526						
vupkhsh	000100	vD	0 0 0 0 0	vB	590						
vupklisb	000100	vD	0 0 0 0 0	vB	654						
vupklsh	000100	vD	0 0 0 0 0	vB	718						
vpkpx	000100	vD	vA	vB	12	782					
vupkhp	000100	vD	0 0 0 0 0	vB	846						
vupklp	000100	vD	0 0 0 0 0	vB	974						
lvsl	011111	vD	A	B	6				0		
lvslr	011111	vD	A	B	38				0		
dst	011111	T	0 0	STRM	A	B	342				0
dstt	011111	1	0 0 0	tag	A	B	0				0
dstst	011111	T	0 0	STRM	A	B	374				0
dststt	011111	1	0 0 0	tag	A	B	11	22	0		
dss	011111	A	0 0	STRM	0 0 0 0 0	0 0 0 0 0	822				0
dssall	011111	A	0 0	STRM	0 0 0 0 0	0 0 0 0 0	822				0
lvebx	011111	vD		A	B	7				0	
lvehx	011111	vD		A	B	39				0	
lviewx	011111	vD		A	B	71				0	
lvx	011111	vD		A	B	103				0	
lvxl	011111	vD		A	B	359				0	

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
stvebx	011111	vS					A					B					135					0						
stvehx	011111	vS					A					B					167					0						
stviewx	011111	vS					A					B					199					0						
stvx	011111	vS					A					B					231					0						
stvxl	011111	vS					A					B					487					0						

PRELIMINARY

## A.3 Instructions Sorted by Form

Table A-3 through Table C-2 list the AltiVec instructions grouped by form.

Key:



Reserved bits

**Table A-3. Loads and Stores (opcode 31)**

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
lvsl	31	vD				A				B				6						0								
lvslr	31	vD				A				B				38						0								
dst	31	T	00		STRM		A				B				342						0							
dstt	31	1	000		tag		A				B				0						0							
dstst	31	T	00		STRM		A				B				374						0							
dststt	31	1	000		tag		A				B				11				22				0					
dss	31	A	00		STRM		00000				00000				822						0							
dssall	31	A	00		STRM		00000				00000				822						0							
lvebx	31	vD				A				B				7						0								
lvehx	31	vD				A				B				39						0								
lvewx	31	vD				A				B				71						0								
lvx	31	vD				A				B				103						0								
lvxl	31	vD				A				B				359						0								
stvebx	31	vS				A				B				135						0								
stvehx	31	vS				A				B				167						0								
stvewx	31	vS				A				B				199						0								
stvx	31	vS				A				B				231						0								
stvxl	31	vS				A				B				487						0								

**Table C-1. VA-Form**

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
vmhaddshs	04	vD				vA				vB				vC				32										
vmhraddshs	04	vD				vA				vB				vC				33										
vmladduhm	04	vD				vA				vB				vC				34										
vmsumubm	04	vD				vA				vB				vC				36										
vmsummbm	04	vD				vA				vB				vC				37										
vmsumuhm	04	vD				A				vB				vC				38										
vmsumuhs	04	vD				vA				vB				vC				39										
vmsumshm	04	vD				vA				vB				vC				40										
vmsumshs	04	vD				vA				vB				vC				41										
vsel	04	vD				vA				vB				vC				42										

**Table C-1. VA-Form (Continued)**

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31			
<b>vperm</b>	04					vD						vA						vB						vC						43	
<b>vsldoi</b>	04					vD						vA						vB		0			SH						44		
<b>vmaddfp</b>	04					vD						vA						vB												46	
<b>vnmsubfp</b>	04					vD						vA						vB						vC						47	

**Table C-2. VX-Form**

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
<b>vaddubm</b>	04				vD					vA				vB														0
<b>vadduhm</b>	04				vD					vA				vB														64
<b>vadduwm</b>	04				vD					vA				vB														128
<b>vaddcuw</b>	04				vD					vA				vB														384
<b>vaddubs</b>	04				vD					vA				vB														512
<b>vadduhs</b>	04				vD					vA				vB														576
<b>vadduws</b>	04				vD					vA				vB														640
<b>vaddsbs</b>	04				vD					vA				vB														768
<b>vaddshs</b>	04				vD					vA				vB														832
<b>vaddsws</b>	04				vD					vA				vB														896
<b>vsububm</b>	04				vD					vA				vB														1024
<b>vsubuhm</b>	04				vD					vA				vB														1088
<b>vsubuwm</b>	04				vD					vA				vB														1152
<b>vsubcuw</b>	04				vD					vA				vB														1408
<b>vsububs</b>	04				vD					vA				vB														1536
<b>vsubuhs</b>	04				vD					vA				vB														1600
<b>vsubuws</b>	04				vD					vA				vB														1664
<b>vsubsbs</b>	04				vD					vA				vB														1792
<b>vsubshs</b>	04				vD					vA				vB														1856
<b>vsubsws</b>	04				vD					vA				vB														1920
<b>vmaxub</b>	04				vD					vA				vB														2
<b>vmaxuh</b>	04				vD					vA				vB														66
<b>vmaxuw</b>	04				vD					vA				vB														130
<b>vmaxsb</b>	04				vD					vA				vB														258
<b>vmaxsh</b>	04				vD					vA				vB														322
<b>vmaxsw</b>	04				vD					vA				vB														386
<b>vminub</b>	04				vD					vA				vB														514
<b>vminuh</b>	04				vD					vA				vB														578
<b>vminuw</b>	04				vD					vA				vB														642

### Table C-2. VX-Form (Continued)

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	20	21	22	23	24	25	26	27	28	29	30	31
vminsb	04					vD				vA					vB												770
vminsh	04					vD				vA					vB												834
vminsw	04					vD				vA					vB												898
vavgub	04					vD				vA					vB												1026
vavguh	04					vD				vA					vB												1090
vavguw	04					vD				vA					vB												1154
vavgusb	04					vD				vA					vB												1282
vavgsh	04					vD				vA					vB												1346
vavgsw	04					vD				vA					vB												1410
vrlb	04					vD				vA					vB												4
vrlh	04					vD				vA					vB												68
vrlw	04					vD				vA					vB												132
vslb	04					vD				vA					vB												260
vslh	04					vD				vA					vB												324
vslw	04					vD				vA					vB												388
vsl	04					vD				vA					vB												452
vsrb	04					vD				vA					vB												516
vsrh	04					vD				vA					vB												580
vsrw	04					vD				vA					vB												644
vsr	04					vD				vA					vB												708
vsrab	04					vD				vA					vB												772
vsrah	04					vD				vA					vB												836
vsraw	04					vD				vA					vB												900
vand	04					vD				vA					vB												1028
vandc	04					vD				vA					vB												1092
vor	04					vD				vA					vB												1156
vxor	04					vD				vA					vB												1220
vnor	04					vD				vA					vB												1284
mfvscr	04					vD				0	0	0	0	0	0	0										1540	0
mtvscr	04					0	0	0	0	0	0	0	0		vB											1604	0
vcmpequbx	04					vD				vA					vB		Rc									6	
vcmpequhx	04					vD				vA					vB		Rc									70	
vcmpequwx	04					vD				vA					vB		Rc									134	
vcmpeqfp	04					vD				vA					vB		Rc									198	
vcmpgefp	04					vD				vA					vB		Rc									454	
vcmpgtubx	04					vD				vA					vB		Rc									518	

**Table C-2. VX-Form (Continued)**

Name	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
<b>vcmpgtuhx</b>	04				vD				vA				vB		Rc								582					
<b>vcmpgtuwX</b>	04				vD				vA				vB		Rc								646					
<b>vcmpgtfpX</b>	04				vD				vA				vB		Rc								710					
<b>vcmpgtsbX</b>	04				vD				vA				vB		Rc								774					
<b>vcmpgtshX</b>	04				vD				vA				vB		Rc								838					
<b>vcmpgtswX</b>	04				vD				vA				vB		Rc								902					
<b>vcmpbfpX</b>	04				vD				vA				vB		Rc								966					
<b>vmuloub</b>	04				vD				vA				vB										8					
<b>vmulouh</b>	04				vD				vA				vB										72					
<b>vmulosb</b>	04				vD				vA				vB										264					
<b>vmulosh</b>	04				vD				vA				vB										328					
<b>vmuleub</b>	04				vD				vA				vB										520					
<b>vmuleuh</b>	04				vD				vA				vB										584					
<b>vmulesb</b>	04				vD				vA				vB										776					
<b>vmulesh</b>	04				vD				vA				vB										840					
<b>vsum4ubs</b>	04				vD				vA				vB										1544					
<b>vsum4sbs</b>	04				vD				vA				vB										1800					
<b>vsum4shs</b>	04				vD				vA				vB										1608					
<b>vsum2sws</b>	04				vD				vA				vB										1672					
<b>vsumsws</b>	04				vD				vA				vB										1928					
<b>vaddfp</b>	04				vD				vA				vB										10					
<b>vsubfp</b>	04				vD				vA				vB										74					
<b>vrrefp</b>	04				vD				0	0	0	0	0	vB									266					
<b>vrsqrtefp</b>	04				vD				0	0	0	0	0	vB									330					
<b>vexptefp</b>	04				vD				0	0	0	0	0	vB									458					
<b>vlogefp</b>	04				vD				0	0	0	0	0	vB									458					
<b>vrfin</b>	04				vD				0	0	0	0	0	vB									522					
<b>vrfiz</b>	04				vD				0	0	0	0	0	vB									586					
<b>vrfip</b>	04				vD				0	0	0	0	0	vB									650					
<b>vrfim</b>	04				vD				0	0	0	0	0	vB									714					
<b>vcfux</b>	04				vD				U	I	M	M	vB										778					
<b>vcfsx</b>	04				vD				U	I	M	M	vB										842					
<b>vctuxs</b>	04				vD				U	I	M	M	vB										906					
<b>vctsxs</b>	04				vD				U	I	M	M	vB										970					
<b>vmaxfp</b>	04				vD				vA				vB										1034					
<b>vminfp</b>	04				vD				vA				vB										1098					

### Table C-2. VX-Form (Continued)

	0	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
vmrghb	04					vD					vA				vB								12					
vmrghh	04					vD					vA				vB								76					
vmrghw	04					vD					vA				vB								140					
vmrglb	04					vD					vA				vB								268					
vmrglh	04					vD					vA				vB								332					
vmrglw	04					vD					vA				vB								396					
vspltb	04					vD					UIMM				vB								524					
vsplth	04					vD					UIMM				vB								588					
vspltw	04					vD					UIMM				vB								652					
vspltisb	04					vD					SIMM			0	0	0	0	0					780					
vspltish	04					vD					SIMM			0	0	0	0	0					844					
vspltisw	04					vD					SIMM			0	0	0	0	0					908					
vslo	04					vD					vA				vB								1036					
vsro	04					vD					vA				vB								1100					
vpkuhum	04					vD					vA				vB								14					
vpkuwum	04					vD					vA				vB								78					
vpkuhus	04					vD					vA				vB								142					
vpkuwus	04					vD					vA				vB								206					
vpkshus	04					vD					vA				vB								270					
vpkswus	04					vD					vA				vB								334					
vpkshss	04					vD					vA				vB								398					
vpkswss	04					vD					vA				vB								462					
vupkhsb	04					vD					0	0	0	0	0	vB							526					
vupkhsh	04					vD					0	0	0	0	0	vB							590					
vupklisb	04					vD					0	0	0	0	0	vB							654					
vupklsh	04					vD					0	0	0	0	0	vB							718					
vpkpx	04					vD					vA				vB			12					782					
vupkhpix	04					vD					0	0	0	0	0	vB							846					
vupklpx	04					vD					0	0	0	0	0	vB							974					

## A.4 Instruction Set Legend

Table A-4 provides general information on the AltiVec instruction set such as the architectural level, privilege level, and form.

**Table A-4. AltiVec Instruction Set Legend**

	UISA	VEA	OEA	Supervisor Level	Optional	Form
<b>dst</b>	√					VX
<b>dstst</b>		√				VX
<b>dststt</b>		√				VX
<b>dstt</b>		√				VX
<b>dss</b>		√				VX
<b>dssall</b>		√				VX
<b>lvebx</b>	√					X
<b>lvehx</b>	√					X
<b>lviewx</b>	√					X
<b>lvsl</b>	√					X
<b>lvslr</b>	√					X
<b>lvx</b>	√					X
<b>lvxl</b>	√					X
<b>mfvscr</b>	√					VX
<b>mtvscr</b>	√					VX
<b>stvebx</b>	√					X
<b>stvehx</b>	√					X
<b>stviewx</b>	√					X
<b>stvx</b>	√					X
<b>stvxl</b>	√					X
<b>vaddcuw</b>	√					VX
<b>vaddfp</b>	√					VX
<b>vaddsbs</b>	√					VX



**Table A-4. AltiVec Instruction Set Legend (Continued)**

	UISA	VEA	OEA	Supervisor Level	Optional	Form
vaddshs	√					VX
vaddsws	√					VX
vaddubm	√					VX
vaddubs	√					VX
vadduhm	√					VX
vadduhs	√					VX
vadduwm	√					VX
vadduws	√					VX
vand	√					VX
vandc	√					VX
vavgsb	√					VX
vavgsh	√					VX
vavgsw	√					VX
vavgub	√					VX
vavguh	√					VX
vavguw	√					VX
vcfux	√					VX
vcfsx	√					VX
vcmpbfp	√					VXR
vcmpqfx	√					VXR
vcmpqubx	√					VXR
vcmpquhx	√					VXR
vcmpquwx	√					VXR
vcmpgefpx	√					VXR
vcmpgtfpx	√					VXR
vcmpgtsbx	√					VXR
vcmpgtshx	√					VXR
vcmpgtswx	√					VXR
vcmpgtubx	√					VXR
vcmpgtuhx	√					VXR
vcmpgtuwx	√					VXR
vctsx	√					VX

**Table A-4. AltiVec Instruction Set Legend (Continued)**

	UISA	VEA	OEA	Supervisor Level	Optional	Form
vctuxs	√					VX
vexptefp	√					VX
vlogefp	√					VX
vmaddfp	√					VA
vmaxfp	√					VX
vmaxsb	√					VX
vmaxsh	√					VX
vmaxsw	√					VX
vmaxub	√					VX
vmaxuh	√					VX
vmaxuw	√					VX
vmhaddshs	√					VA
vmhraddshs	√					VA
vminfp	√					VX
vminsb	√					VX
vminsh	√					VX
vminsw	√					VX
vminub	√					VX
vminuh	√					VX
vminuw	√					VX
vmladduhm	√					VA
vmrghb	√					VX
vmrghh	√					VX
vmrghw	√					VX
vmrglb	√					VX
vmrglh	√					VX
vmrglw	√					VX
vmsummbm	√					VA
vmsumshm	√					VA
vmsumshs	√					VA
vmsumubm	√					VA
vmsumuhm	√					VA

**Table A-4. AltiVec Instruction Set Legend (Continued)**

	UISA	VEA	OEA	Supervisor Level	Optional	Form
<b>vmsumuhs</b>	√					VA
<b>vmulesb</b>	√					VX
<b>vmulesh</b>	√					VX
<b>vmuleub</b>	√					VX
<b>vmuleuh</b>	√					VX
<b>vmulosb</b>	√					VX
<b>vmulosh</b>	√					VX
<b>vmuloub</b>	√					VX
<b>vmulouh</b>	√					VX
<b>vnmsubfp</b>	√					VA
<b>vnor</b>	√					VX
<b>vor</b>	√					VX
<b>vperm</b>	√					VA
<b>vpkpx</b>	√					VX
<b>vpkshss</b>	√					VX
<b>vpkshus</b>	√					VX
<b>vpkswss</b>	√					VX
<b>vpkuhum</b>	√					VX
<b>vpkuhus</b>	√					VX
<b>vpkswus</b>	√					VX
<b>vpkuwum</b>	√					VX
<b>vpkuwus</b>	√					VX
<b>vrefp</b>	√					VX
<b>vrfim</b>	√					VX
<b>vrfim</b>	√					VX
<b>vrfin</b>	√					VX
<b>vrfip</b>	√					VX
<b>vrfiz</b>	√					VX
<b>vrlb</b>	√					VX
<b>vrlh</b>	√					VX
<b>vrlw</b>	√					VX
<b>vrsqrtefp</b>	√					VX
<b>vsel</b>	√					VA

**Table A-4. AltiVec Instruction Set Legend (Continued)**

	UISA	VEA	OEA	Supervisor Level	Optional	Form
<b>vsl</b>	√					VX
<b>vsib</b>	√					VX
<b>vsldoi</b>	√					VA
<b>vslih</b>	√					VX
<b>vslo</b>	√					VX
<b>vslw</b>	√					VX
<b>vspltb</b>	√					VX
<b>vsplth</b>	√					VX
<b>vspltisb</b>	√					VX
<b>vspltish</b>	√					VX
<b>vspltisw</b>	√					VX
<b>vspltw</b>	√					VX
<b>vsr</b>	√					VX
<b>vsrab</b>	√					VX
<b>vsrah</b>	√					VX
<b>vsraw</b>	√					VX
<b>vsrb</b>	√					VX
<b>vsrh</b>	√					VX
<b>vsro</b>	√					VX
<b>vsrw</b>	√					VX
<b>vsubcuw</b>	√					VX
<b>vsubfp</b>	√					VX
<b>vsubsbs</b>	√					VX
<b>vsubshs</b>	√					VX
<b>vsubsws</b>	√					VX
<b>vsububm</b>	√					VX
<b>vsubuhm</b>	√					VX
<b>vsububs</b>	√					VX
<b>vsubuhs</b>	√					VX
<b>vsubuwm</b>	√					VX
<b>vsubuws</b>	√					VX
<b>vsumsws</b>	√					VX

**Table A-4. AltiVec Instruction Set Legend (Continued)**

	UISA	VEA	OEA	Supervisor Level	Optional	Form
<b>vsum2sws</b>	√					VX
<b>vsum4sbs</b>	√					VX
<b>vsum4shs</b>	√					VX
<b>vsum4ubs</b>	√					VX
<b>vupkhpX</b>	√					VX
<b>vupkhsb</b>	√					VX
<b>vupklsh</b>	√					VX
<b>vupkhpX</b>	√					VX
<b>vupklsb</b>	√					VX
<b>vupklsh</b>	√					VX
<b>vxor</b>	√					VX

**PRELIMINARY**

**PRELIMINARY**

# Glossary of Terms and Abbreviations

A((add persistence and transience to glossary—also swizzle and munge, stride, modulo addressing)))

The glossary contains an alphabetical list of terms, phrases, and abbreviations used in this book. Some of the terms and definitions included in the glossary are reprinted from *IEEE Std. 754-1985, IEEE Standard for Binary Floating-Point Arithmetic*, copyright ©1985 by the Institute of Electrical and Electronics Engineers, Inc. with the permission of the IEEE.

Note that some terms are defined in the context of how they are used in this book.

---

## A

**Architecture.** A detailed specification of requirements for a processor or computer system. It does not specify details of how the processor or computer system must be implemented; instead it provides a template for a family of compatible *implementations*.

**Asynchronous exception.** *Exceptions* that are caused by events external to the processor's execution. In this document, the term 'asynchronous exception' is used interchangeably with the word *interrupt*.

**Atomic access.** A bus access that attempts to be part of a read-write operation to the same address uninterrupted by any other access to that address (the term refers to the fact that the transactions are indivisible). The PowerPC architecture implements atomic accesses through the **lwarx/stwcx** instruction pair.

---

## B

**BAT (block address translation) mechanism.** A software-controlled array that stores the available block address translations on-chip.

**Biased exponent.** An *exponent* whose range of values is shifted by a constant (bias). Typically a bias is provided to allow a range of positive values to express a range that includes both positive and negative values.

**Big-endian.** A byte-ordering method in memory where the address *n* of a word corresponds to the *most-significant byte*. In an addressed

memory word, the bytes are ordered (left to right) 0, 1, 2, 3, with 0 being the most-significant byte. *See* Little-endian.

**Block.** An area of memory that ranges from 128 Kbyte to 256 Mbyte, whose size, translation, and protection attributes are controlled by the *BAT mechanism*.

**Boundedly undefined.** A characteristic of results of certain operations that are not rigidly prescribed by the PowerPC architecture. Boundedly-undefined results for a given operation may vary among implementations, and between execution attempts in the same implementation.

Although the architecture does not prescribe the exact behavior for when results are allowed to be boundedly undefined, the results of executing instructions in contexts where results are allowed to be boundedly undefined are constrained to ones that could have been achieved by executing an arbitrary sequence of defined instructions, in valid form, starting in the state the machine was in before attempting to execute the given instruction.

---

## C

**Cache.** High-speed memory component containing recently-accessed data and/or instructions (subset of main memory).

**Cache block.** A small region of contiguous memory that is copied from memory into a *cache*. The size of a cache block may vary among processors; the maximum block size is one *page*. In PowerPC processors, *cache coherency* is maintained on a cache-block basis. Note that the term ‘cache block’ is often used interchangeably with ‘cache line’.

**Cache coherency.** An attribute wherein an accurate and common view of memory is provided to all devices that share the same memory system. Caches are coherent if a processor performing a read from its cache is supplied with data corresponding to the most recent value written to memory or to another processor’s cache.

**Cache flush.** An operation that removes from a cache any data from a specified address range. This operation ensures that any modified data within the specified address range is written back to main memory. This operation is generated typically by a Data Cache Block Flush (**dcbf**) instruction.

**Caching-inhibited.** A memory update policy in which the *cache* is bypassed and the load or store is performed to or from main memory.



**Cast-outs.** *Cache blocks* that must be written to memory when a cache miss causes a cache block to be replaced.

**Changed bit.** One of two *page history bits* found in each *page table entry* (PTE). The processor sets the changed bit if any store is performed into the *page*. *See also* Page access history bits and Referenced bit.

**Clear.** To cause a bit or bit field to register a value of zero. *See also* Set.

**Context synchronization.** An operation that ensures that all instructions in execution complete past the point where they can produce an *exception*, that all instructions in execution complete in the context in which they began execution, and that all subsequent instructions are *fetches* and executed in the new context. Context synchronization may result from executing specific instructions (such as **isync** or **rfi**) or when certain events occur (such as an exception).

**Copy-back.** An operation in which modified data in a *cache block* is copied back to memory.

---

## D

**Denormalized number.** A nonzero floating-point number whose *exponent* has a reserved value, usually the format's minimum, and whose explicit or implicit leading significand bit is zero.

**Direct-mapped cache.** A cache in which each main memory address can appear in only one location within the cache, operates more quickly when the memory request is a cache hit.

**Direct-store.** Interface available on PowerPC processors only to support direct-store devices from the POWER architecture. When the T bit of a *segment descriptor* is set, the descriptor defines the region of memory that is to be used as a direct-store segment. Note that this facility is being phased out of the architecture and will not likely be supported in future devices. Therefore, software should not depend on it and new software should not use it.

---

## E

**Effective address (EA).** The 32- or 64-bit address specified for a load, store, or an instruction fetch. This address is then submitted to the MMU for translation to either a *physical memory* address or an I/O address.

**Exception.** A condition encountered by the processor that requires special, supervisor-level processing.

**Exception handler.** A software routine that executes when an exception is taken. Normally, the exception handler corrects the condition that

caused the exception, or performs some other meaningful task (that may include aborting the program that caused the exception). The address for each exception handler is identified by an exception vector offset defined by the architecture and a prefix selected via the MSR.

**Extended opcode.** A secondary opcode field generally located in instruction bits 21–30, that further defines the instruction type. All PowerPC instructions are one word in length. The most significant 6 bits of the instruction are the *primary opcode*, identifying the type of instruction. *See also* Primary opcode.

**Execution synchronization.** A mechanism by which all instructions in execution are architecturally complete before beginning execution (appearing to begin execution) of the next instruction. Similar to context synchronization but doesn't force the contents of the instruction buffers to be deleted and refetched.

**Exponent.** In the binary representation of a floating-point number, the exponent is the component that normally signifies the integer power to which the value two is raised in determining the value of the represented number. *See also* Biased exponent.

---

## F

**Fetch.** Retrieving instructions from either the cache or main memory and placing them into the instruction queue.

**Floating-point register (FPR).** Any of the 32 registers in the floating-point register file. These registers provide the source operands and destination results for floating-point instructions. Load instructions move data from memory to FPRs and store instructions move data from FPRs to memory. The FPRs are 64 bits wide and store floating-point values in double-precision format.

**Fraction.** In the binary representation of a floating-point number, the field of the *significand* that lies to the right of its implied binary point.

**Fully-associative.** Addressing scheme where every cache location (every byte) can have any possible address.

---

## G

**General-purpose register (GPR).** Any of the 32 registers in the general-purpose register file. These registers provide the source operands and destination results for all integer data manipulation instructions. Integer load instructions move data from memory to GPRs and store instructions move data from GPRs to memory.

**Guarded.** The guarded attribute pertains to out-of-order execution. When a page is designated as guarded, instructions and data cannot be accessed out-of-order.

---

## H

**Harvard architecture.** An architectural model featuring separate caches for instruction and data.

**Hashing.** An algorithm used in the *page table* search process.

---

## I

**IEEE 754.** A standard written by the Institute of Electrical and Electronics Engineers that defines operations and representations of binary floating-point arithmetic.

**Illegal instructions.** A class of instructions that are not implemented for a particular PowerPC processor. These include instructions not defined by the PowerPC architecture. In addition, for 32-bit implementations, instructions that are defined only for 64-bit implementations are considered to be illegal instructions. For 64-bit implementations instructions that are defined only for 32-bit implementations are considered to be illegal instructions.

**Implementation.** A particular processor that conforms to the PowerPC architecture, but may differ from other architecture-compliant implementations for example in design, feature set, and implementation of *optional* features. The PowerPC architecture has many different implementations.

**Implementation-dependent.** An aspect of a feature in a processor's design that is defined by a processor's design specifications rather than by the PowerPC architecture.

**Implementation-specific.** An aspect of a feature in a processor's design that is not required by the PowerPC architecture, but for which the PowerPC architecture may provide concessions to ensure that processors that implement the feature do so consistently.

**Imprecise exception.** A type of *synchronous exception* that is allowed not to adhere to the precise exception model (*see* Precise exception). The PowerPC architecture allows only floating-point exceptions to be handled imprecisely.

**Inexact.** Loss of accuracy in an arithmetic operation when the rounded result differs from the infinitely precise value with unbounded range.

**In-order.** An aspect of an operation that adheres to a sequential model. An operation is said to be performed in-order if, at the time that it is performed, it is known to be required by the sequential execution model. *See* Out-of-order.

**Instruction latency.** The total number of clock cycles necessary to execute an instruction and make ready the results of that instruction.

**Instruction parallelism.** A feature of PowerPC processors that allows instructions to be processed in parallel.

**Interrupt.** An *asynchronous exception*. On PowerPC processors, interrupts are a special case of exceptions. *See also* asynchronous exception.

**Invalid state.** State of a cache entry that does not currently contain a valid copy of a cache block from memory.

---

## K

**Key bits.** A set of key bits referred to as Ks and Kp in each segment register and each BAT register. The key bits determine whether supervisor or user programs can access a *page* within that *segment* or *block*.

**Kill.** An operation that causes a *cache block* to be invalidated.

---

## L

**L2 cache.** *See* Secondary cache.

**Least-significant bit (lsb).** The bit of least value in an address, register, data element, or instruction encoding.

**Least-significant byte (LSB).** The byte of least value in an address, register, data element, or instruction encoding.

**Little-endian.** A byte-ordering method in memory where the address *n* of a word corresponds to the *least-significant byte*. In an addressed memory word, the bytes are ordered (left to right) 3, 2, 1, 0, with 3 being the *most-significant byte*. *See* Big-endian.

---

## M

**MESI (modified/exclusive/shared/invalid).** *Cache coherency* protocol used to manage caches on different devices that share a memory system. Note that the PowerPC architecture does not specify the implementation of a MESI protocol to ensure cache coherency.

**Memory access ordering.** The specific order in which the processor performs load and store memory accesses and the order in which those accesses complete.

**Memory-mapped accesses.** Accesses whose addresses use the page or block address translation mechanisms provided by the MMU and that occur externally with the bus protocol defined for memory.

**Memory coherency.** An aspect of caching in which it is ensured that an accurate view of memory is provided to all devices that share system memory.

**Memory consistency.** Refers to agreement of levels of memory with respect to a single processor and system memory (for example, on-chip cache, secondary cache, and system memory).

**Memory management unit (MMU).** The functional unit that is capable of translating an *effective* (logical) *address* to a physical address, providing protection mechanisms, and defining caching methods.

**Microarchitecture.** The hardware details of a microprocessor's design. Such details are not defined by the PowerPC architecture.

**Mnemonic.** The abbreviated name of an instruction used for coding.

**Modified state.** When a cache block is in the modified state, it has been modified by the processor since it was copied from memory. *See* MESI.

**Munging.** A modification performed on an *effective address* that allows it to appear to the processor that individual aligned scalars are stored as *little-endian* values, when in fact it is stored in *big-endian* order, but at different byte addresses within double words. Note that munging affects only the effective address and not the byte order. Note also that this term is not used by the PowerPC architecture.

**Multiprocessing.** The capability of software, especially operating systems, to support execution on more than one processor at the same time.

**Most-significant bit (msb).** The highest-order bit in an address, registers, data element, or instruction encoding.

**Most-significant byte (MSB).** The highest-order byte in an address, registers, data element, or instruction encoding.

---

## N

**NaN.** An abbreviation for 'Not a Number'; a symbolic entity encoded in floating-point format. There are two types of NaNs—signaling NaNs (SNaNs) and quiet NaNs (QNaNs).

**No-op.** No-operation. A single-cycle operation that does not affect registers or generate bus activity.

**Normalization.** A process by which a floating-point value is manipulated such that it can be represented in the format for the appropriate precision (single- or double-precision). For a floating-point value to be representable in the single- or double-precision format, the leading implied bit must be a 1.

---

## O

**OEA (operating environment architecture).** The level of the architecture that describes PowerPC memory management model, supervisor-level registers, synchronization requirements, and the exception model. It also defines the time-base feature from a supervisor-level perspective. Implementations that conform to the PowerPC OEA also conform to the PowerPC UISA and VEA.

**Optional.** A feature, such as an instruction, a register, or an exception, that is defined by the PowerPC architecture but not required to be implemented.

**Out-of-order.** An aspect of an operation that allows it to be performed ahead of one that may have preceded it in the sequential model, for example, speculative operations. An operation is said to be performed out-of-order if, at the time that it is performed, it is not known to be required by the sequential execution model. *See* In-order.

**Out-of-order execution.** A technique that allows instructions to be issued and completed in an order that differs from their sequence in the instruction stream.

**Overflow.** An error condition that occurs during arithmetic operations when the result cannot be stored accurately in the destination register(s). For example, if two 32-bit numbers are multiplied, the result may not be representable in 32 bits.

---

## P

**Page.** A region in memory. The OEA defines a page as a 4-Kbyte area of memory, aligned on a 4-Kbyte boundary.

**Page access history bits.** The *changed* and *referenced* bits in the PTE keep track of the access history within the page. The referenced bit is set by the MMU whenever the page is accessed for a read or write operation. The changed bit is set when the page is stored into. *See* Changed bit and Referenced bit.

**Page fault.** A page fault is a condition that occurs when the processor attempts to access a memory location that does not reside within a *page* not currently resident in *physical memory*. On PowerPC

processors, a page fault exception condition occurs when a matching, valid *page table entry* (PTE[V] = 1) cannot be located.

**Page table.** A table in memory is comprised of *page table entries*, or PTEs. It is further organized into eight PTEs per PTEG (page table entry group). The number of PTEGs in the page table depends on the size of the page table (as specified in the SDR1 register).

**Page table entry (PTE).** Data structures containing information used to translate *effective address* to physical address on a 4-Kbyte page basis. A PTE consists of 8 bytes of information in a 32-bit processor and 16 bytes of information in a 64-bit processor.

**Physical memory.** The actual memory that can be accessed through the system's memory bus.

**Pipelining.** A technique that breaks operations, such as instruction processing or bus transactions, into smaller distinct stages or tenures (respectively) so that a subsequent operation can begin before the previous one has completed.

**Precise exceptions.** A category of exception for which the pipeline can be stopped so instructions that preceded the faulting instruction can complete, and subsequent instructions can be flushed and redispached after exception handling has completed. *See* Imprecise exceptions.

**Primary opcode.** The most-significant 6 bits (bits 0–5) of the instruction encoding that identifies the type of instruction. *See* Secondary opcode.

**Protection boundary.** A boundary between *protection domains*.

**Protection domain.** A protection domain is a segment, a virtual page, a BAT area, or a range of unmapped effective addresses. It is defined only when the appropriate relocate bit in the MSR (IR or DR) is 1.

---

## Q

**Quad word.** A group of 16 contiguous locations starting at an address divisible by 16.

**Quiet NaN.** A type of *NaN* that can propagate through most arithmetic operations without signaling exceptions. A quiet NaN is used to represent the results of certain invalid operations, such as invalid arithmetic operations on infinities or on NaNs, when invalid. *See* Signaling NaN.

## R

**rA.** The **rA** instruction field is used to specify a GPR to be used as a source or destination.

**rB.** The **rB** instruction field is used to specify a GPR to be used as a source.

**rD.** The **rD** instruction field is used to specify a GPR to be used as a destination.

**rS.** The **rS** instruction field is used to specify a GPR to be used as a source.

**Real address mode.** An MMU mode when no address translation is performed and the *effective address* specified is the same as the physical address. The processor's MMU is operating in real address mode if its ability to perform address translation has been disabled through the MSR registers IR and/or DR bits.

**Record bit.** Bit 31 (or the Rc bit) in the instruction encoding. When it is set, updates the condition register (CR) to reflect the result of the operation.

**Referenced bit.** One of two *page history bits* found in each *page table entry* (PTE). The processor sets the *referenced bit* whenever the page is accessed for a read or write. *See also* Page access history bits.

**Register indirect addressing.** A form of addressing that specifies one GPR that contains the address for the load or store.

**Register indirect with immediate index addressing.** A form of addressing that specifies an immediate value to be added to the contents of a specified GPR to form the target address for the load or store.

**Register indirect with index addressing.** A form of addressing that specifies that the contents of two GPRs be added together to yield the target address for the load or store.

**Reservation.** The processor establishes a reservation on a *cache block* of memory space when it executes an **lwarx** instruction to read a memory semaphore into a GPR.

**Reserved field.** In a register, a reserved field is one that is not assigned a function. A reserved field may be a single bit. The handling of reserved bits is *implementation-dependent*. Software is permitted to write any value to such a bit. A subsequent reading of the bit returns 0 if the value last written to the bit was 0 and returns an undefined value (0 or 1) otherwise.



**RISC (reduced instruction set computing).** An *architecture* characterized by fixed-length instructions with nonoverlapping functionality and by a separate set of load and store instructions that perform memory accesses.

---

## S

**Scalability.** The capability of an architecture to generate *implementations* specific for a wide range of purposes, and in particular implementations of significantly greater performance and/or functionality than at present, while maintaining compatibility with current implementations.

**Secondary cache.** A cache memory that is typically larger and has a longer access time than the primary cache. A secondary cache may be shared by multiple devices. Also referred to as L2, or level-2, cache.

**Segment.** A 256-Mbyte area of *virtual memory* that is the most basic memory space defined by the PowerPC architecture. Each segment is configured through a unique *segment descriptor*.

**Segment descriptors.** Information used to generate the interim *virtual address*. The segment descriptors reside in 16 on-chip segment registers for 32-bit implementations. For 64-bit implementations, the segment descriptors reside as *segment table entries* in a hashed segment table in memory.

**Set (v).** To write a nonzero value to a bit or bit field; the opposite of *clear*. The term ‘set’ may also be used to generally describe the updating of a bit or bit field.

**Set (n).** A subdivision of a *cache*. Cacheable data can be stored in a given location in any one of the sets, typically corresponding to its lower-order address bits. Because several memory locations can map to the same location, cached data is typically placed in the set whose *cache block* corresponding to that address was used least recently. *See* Set-associative.

**Set-associative.** Aspect of cache organization in which the cache space is divided into sections, called *sets*. The cache controller associates a particular main memory address with the contents of a particular set, or region, within the cache.

**Signaling NaN.** A type of *NaN* that generates an invalid operation program exception when it is specified as arithmetic operands. *See* Quiet NaN.

**Significand.** The component of a binary floating-point number that consists of an explicit or implicit leading bit to the left of its implied binary point and a fraction field to the right.

**Simplified mnemonics.** Assembler mnemonics that represent a more complex form of a common operation.

**Splat.** A splat instruction will take one element and replicate (splat) that value into a vector register. The purpose being to have all elements have the same value so they can be used as a constant to multiply other vector registers.

**Static branch prediction.** Mechanism by which software (for example, compilers) can give a hint to the machine hardware about the direction a branch is likely to take.

**Sticky bit.** A bit that when *set* must be cleared explicitly.

**Strong ordering.** A memory access model that requires exclusive access to an address before making an update, to prevent another device from using stale data.

**Superscalar machine.** A machine that can issue multiple instructions concurrently from a conventional linear instruction stream.

**Supervisor mode.** The privileged operation state of a processor. In supervisor mode, software, typically the operating system, can access all control registers and can access the supervisor memory space, among other privileged operations.

**Synchronization.** A process to ensure that operations occur strictly *in order*. See Context synchronization and Execution synchronization.

**Synchronous exception.** An *exception* that is generated by the execution of a particular instruction or instruction sequence. There are two types of synchronous exceptions, *precise* and *imprecise*.

**System memory.** The physical memory available to a processor.

---

**T** **TLB (translation lookaside buffer)** A cache that holds recently-used *page table entries*.

**Throughput.** The measure of the number of instructions that are processed per clock cycle.

**Tiny.** A floating-point value that is too small to be represented for a particular precision format, including *denormalized* numbers; they do not include  $\pm 0$ .

---

**U** **UISA (user instruction set architecture).** The level of the architecture to which user-level software should conform. The UISA defines the base user-level instruction set, user-level registers, data types, floating-point memory conventions and exception model as seen by user programs, and the memory and programming models.

**Underflow.** An error condition that occurs during arithmetic operations when the result cannot be represented accurately in the destination register. For example, underflow can happen if two floating-point fractions are multiplied and the result requires a smaller *exponent* and/or mantissa than the single-precision format can provide. In other words, the result is too small to be represented accurately.

**Unified cache.** Combined data and instruction cache.

**User mode.** The unprivileged operating state of a processor used typically by application software. In user mode, software can only access certain control registers and can access only user memory space. No privileged operations can be performed. Also referred to as problem state.

---

**V** **VEA (virtual environment architecture).** The level of the *architecture* that describes the memory model for an environment in which multiple devices can access memory, defines aspects of the cache model, defines cache control instructions, and defines the time-base facility from a user-level perspective. *Implementations* that conform to the PowerPC VEA also adhere to the UISA, but may not necessarily adhere to the OEA.

**Virtual address.** An intermediate address used in the translation of an *effective address* to a physical address.

**Virtual memory.** The address space created using the memory management facilities of the processor. Program access to virtual memory is possible only when it coincides with *physical memory*.

---

## W

**Weak ordering.** A memory access model that allows bus operations to be reordered dynamically, which improves overall performance and in particular reduces the effect of memory latency on instruction throughput.

**Word.** A 32-bit data element.

**Write-back.** A cache memory update policy in which processor write cycles are directly written only to the cache. External memory is updated only indirectly, for example, when a modified cache block is *cast out* to make room for newer data.

**Write-through.** A cache memory update policy in which all processor write cycles are written to both the cache and memory.

# INDEX

## A

- Accesses
  - misaligned accesses, 3-1
- Acronyms and abbreviated terms, list, xxiv
- add, 4-6, 4-6, 4-7, 6-18, 6-20, 6-21
- addi, 4-5, 4-7, 4-7, 4-8, 4-8, 4-8, 4-8, 4-9, 4-9
- addis, 4-6
- Address calculation
  - load and store instructions, 4-24
- Addressing conventions
  - alignment, 3-1
  - byte ordering, 3-2
- Addressing modes
  - register indirect
    - with index, integer, 4-24
- Aligned scalars, LE mode, 3-5
- Alignment
  - rules, 3-1, 3-5
- andi., 4-14, 4-14, 4-15, 4-15, 4-15
- Arithmetic instructions
  - floating-point, 4-18
  - integer, 4-1

## B

- Big-endian mode
  - byte ordering, 1-7, 3-3
  - concept, 3-3
- Byte ordering
  - aligned scalars, LE mode, 3-5
  - big-endian mode, default, 3-2, 3-3
  - concept, 3-2
  - default, 1-7, 4-3
  - LE and ILE bits in MSR, 3-2
  - least-significant byte (LSB), 3-3
  - little-endian mode
    - description, 3-3
  - most-significant byte (MSB), 3-3

## C

- Cache management instructions
  - dcbt, 4-39, 4-39
  - dcbst, 4-40, 4-40
  - list of instructions, 4-39
- Classes of instructions, 4-2, 4-2
- cmp, 4-11, 4-12, 4-12, 4-12, 4-13, 4-14, 4-14, 4-14, 4-18, 4-18, 4-22, 4-22, 4-22, 4-23, 4-34

- Compare instructions
  - floating-point, 4-20
  - integer, 4-12
- Computation modes
  - effective address, 4-2
  - PowerPC architecture, 1-4, 4-2
- Conventions
  - instruction set
    - classes of instructions, 4-2
    - computation modes, 4-2
    - memory addressing, 4-3
  - operand conventions
    - architecture levels represented, 3-1
  - terminology, xxvi
- CR (condition register)
  - bit fields, 2-5
  - CRn field, compare instructions, 2-5
  - move to/from CR instructions, 4-36

## D

- Data organization, memory, 3-1
- Data types
  - aligned scalars, 3-5
- dcbt, 4-39, 4-39
- dcbst, 4-40, 4-40
- dcbz, 4-41, 4-41

## E

- Effective address calculation
  - branches, 4-3
  - EA modifications, 3-5
  - loads and stores, 4-3, 4-24
- Exceptions
  - exception model, overview, 1-11
  - overview, 1-11
- Exclusive OR (XOR), 3-5
- Execution model
  - floating-point, 3-12
- Extended mnemonics, *see* Simplified mnemonics

## F

- fadds, 4-18, 4-18
- fctiw, 4-20, 4-20, 4-20, 4-20
- Floating-point model
  - execution model
    - floating-point, 3-12

# INDEX

- FP arithmetic instructions, 4-18
- FP compare instructions, 4-20
- FP execution model, 3-12
- FP rounding/conversion instructions, 4-19
- fmadds, 4-19, 4-19
- FPSCR (floating-point status and control register)
  - bit settings, 2-3, 2-5
- frsp, 4-20, 4-20, 4-20, 4-20

## I

- Instruction set conventions
  - classes of instructions, 4-2
  - computation modes, 4-2
  - memory addressing, 4-3
- Instructions
  - cache management instructions
    - dcbt, 4-39, 4-39
    - dcbtst, 4-40, 4-40
    - list of instructions, 4-39
  - classes of instructions, 4-2
  - floating-point
    - arithmetic, 4-18
    - compare, 4-20
    - computational instructions, 3-12
    - noncomputational instructions, 3-12
    - rounding/conversion, 4-19
  - instruction field conventions, xxvii
  - integer
    - arithmetic, 4-1, 4-4
    - compare, 4-12
    - load, 4-25, 4-26
    - logical, 4-1, 4-14
    - rotate/shift, 4-15-??
    - store, 4-28
  - load and store
    - address generation, integer, 4-24
    - integer load, 4-25, 4-26
    - integer store, 4-28
  - memory control instructions, 4-38
  - PowerPC instructions, list, A-1, A-18
  - processor control instructions, 4-36
- Integer arithmetic instructions, 4-1, 4-4
- Integer compare instructions, 4-12
- Integer load instructions, 4-25, 4-26
- Integer logical instructions, 4-1, 4-14
- Integer rotate/shift instructions, 4-15-??
- Integer store instructions
  - description, 4-28

## L

- lbzx, 4-26, 4-26, 4-26, 4-27, 4-27, 4-28, 4-28
- Little-endian mode
  - byte ordering, 3-2, 3-3

- description, 3-3
- LE and ILE bits, 3-2
- mapping, 3-4
- Load/store
  - address generation, integer, 4-24
  - integer load instructions, 4-25, 4-26
  - integer store instructions, 4-28
- Logical instructions, integer, 4-1, 4-14

## M

- Memory addressing, 4-3
- Memory control instructions
  - user-level cache, 4-38
- Memory management unit
  - overview, 1-11
- Memory operands, 4-3
- Memory, data organization, 3-1
- mffs, 4-23, 4-23
- Misaligned accesses and alignment, 3-1
- Move to/from CR instructions, 4-36
- MSR (machine state register)
  - bit settings, 2-7
  - LE and ILE bits, 3-2
- mterf, 4-36, 4-37
- Munging
  - description, 3-5

## O

- OEA (operating environment architecture)
  - definition, xx, 1-5
  - programming model, 2-10
  - register set, 2-9
- Operands
  - conventions, description, 1-7, 3-1
  - memory operands, 4-3
- Operating environment architecture, *see* OEA

## P

- PowerPC architecture
  - computation modes, 1-4, 4-2
  - features summary
    - defined features, 1-3
    - features not defined, 1-5
  - instruction list, A-1, A-18
  - levels of the PowerPC architecture, 1-5-??
  - operating environment architecture, xx, 1-5
  - overview, 1-2
  - registers
    - programming model, 1-6, 2-10
  - user instruction set architecture, xix, 1-5
  - virtual environment architecture, xix, 1-5
- Privilege levels

# INDEX

user-level cache control instructions, 4-38  
Processor control instructions, 4-36  
Programming model  
    all registers (OEA), 2-10

## R

Record bit (Rc)  
    description, 6-2  
Registers  
    exception handling registers  
        SRR0/SRR1, 2-8  
    OEA register set, 2-9  
    supervisor-level  
        SRR0/SRR1, 2-8  
    UISA register set, 2-1  
    user-level  
        CR, 2-5  
        XER, 2-4  
Rotate/shift instructions, 4-15–??  
Rounding/conversion instructions, FP, 4-19

## S

Scalars  
    aligned, LE mode, 3-5  
Shift/rotate instructions, 4-15–??  
Simplified mnemonics  
    recommended mnemonics, 4-37  
SRR0/SRR1 (status save/restore registers)  
    format, 2-8, 2-8  
stbx, 4-28, 4-28, 4-28  
subf, 4-6, 4-7, 4-7

## T

Terminology conventions, xxvi

## U

UISA (user instruction set architecture)  
    definition, xix, 1-5  
    register set, 2-1  
User instruction set architecture, *see* UISA

## V

VEA (virtual environment architecture)  
    definition, xix, 1-5  
Virtual environment architecture, *see* VEA

## X

XOR (exclusive OR), 3-5

# INDEX

**PRELIMINARY**



Overview	1
AltiVec Register Set	2
Operand Conventions	3
Addressing Modes and Instruction Set Summary	4
Cache, Exceptions, and Memory Management	5
AltiVec Instructions	6
AltiVec Instruction Set Listings	A
Glossary of Terms and Abbreviations	GLO
Index	IND

1

Overview

2

Altivec Register Set

3

Operand Conventions

4

Addressing Modes and Instruction Set Summary

5

Cache, Exceptions, and Memory Management

6

Altivec Instructions

A

Altivec Instruction Set Listings

GLO

Glossary of Terms and Abbreviations

IND

Index

# Attention!

This book is a companion to the *PowerPC Microprocessor Family: The Programming Environments*, referred to as *The Programming Environments Manual*. Note that the companion *Programming Environments Manual* exists in two versions. See the Preface for a description of the following two versions:

- *PowerPC Microprocessor Family: The Programming Environments*, Rev 1  
Order #: MPCFPE/AD
- *PowerPC Microprocessor Family: The Programming Environments for 32-Bit Microprocessors*, Rev 1  
Order #: MPCFPE32B/AD

Call the Motorola LDC at 1-800-441-2447 (website: <http://ldc.nmd.com>) or contact your local sales office to obtain copies.

